The George Floyd Effect: How Protests and Public Scrutiny Changed Police Behavior

Marcel Roman[1], Klara Fredriksson[2], Chris Cassella[3], Derek Epp[4], and Hannah Walker[5]

[1]Postdoctoral Fellow, Harvard University

[2]Assistant Professor, Beloit College

[3]PhD Student, University of Texas at Austin

[4]Associate Professor, University of Texas at Austin

[5]Associate Professor, University of Texas at Austin

September 13, 2023

Abstract

The murder of George Floyd in May 2020 sparked a wave of Black Lives Matter protests in many cities throughout the United States. Protesters demanded constraints against the police and policing. These have led some to worry about the possibility of a "Ferguson Effect," where police withdraw from policing, and in particular discretionary stops and searches, with deleterious consequences for crime. Drawing on data from four cities, we evaluate whether the 2020 BLM protests impacted police behavior, and whether changes in policing negatively impacted public safety. Regression discontinuity-in-time estimates suggest that although depolicing followed the BLM protests, in some respects the quality of policing improved, and public safety was not clearly impacted. Our findings have important implications for research on policing, social movements, and structural inequality in cities.

Keywords: depolicing; bureaucratic accountability; bureaucratic politics; Black Lives Matter

Introduction

George Floyd was murdered by police on May 25, 2020. Police officers handcuffed him, pinned him to the ground, and officer Derek Chauvin knelt on his neck for almost nine minutes, ending his life. A video of the incident quickly went viral and sparked the most significant wave of Black Lives Matter (BLM) protests to date. By the first week of June, 2020, protests had occurred in over 140 cities across the US and extended to over 40 countries (Smith, Wu, and Murphy, 2020). In addition to their unprecedented scale, the 2020 BLM protests were tonally radical, pushing the language of abolition into the mainstream and redefining the discourse around policing. While calls to *defund the police* proved politically incendiary and the demands of activists varied, a desire to reduce police brutality, reduce overall contact between citizens and police, and hold police accountable for misconduct animated the movement.

Anecdotal accounts across various media outlets suggest that the protests led to a decline in policing (whether because officers were defunded, demoralized, or counter-protesting) and in turn a rise in crime (Arango, 2021; Pagones, 2020). But this is speculation. Little empirical evidence exists connecting depolicing and protests to crime, and no evidence that we know of connects the protests of 2020 specifically to rising violence in cities. As such, the impact of the protests on both policing practices and public safety remains an open empirical question with implications for our understanding of policing, protests, and bureaucratic responsiveness.

Existing literature suggests two reasons why police activity may decline following protests. Police may respond to the demands of protesters by changing tactics in ways that reduce contact with citizens, improve the efficiency of their work, and, especially, reduce racially unequal outcomes. Thus, depolicing may be a response to protester demands. Researchers elsewhere demonstrate that protests do have the ability to hold public officials accountable (Gillion, 2012; Gause, 2022). In contrast, police may change their tactics in ways springing from demoralization, burnout, or retaliation against external critiques. Scholars have referred to this kind of behavior as dissent shirking (Chanin and Sheats, 2018). While a handful of studies tackle whether depolicing occurs, very few characterize withdrawal, raising questions around the reasons behind this behavior and its consequences for civilians (Nix, Wolfe, and Campbell, 2018).

To address these questions, we evaluate several years of high resolution, incident level data in four contexts: Seattle, WA, Austin, TX, Philadelphia, PA and Los Angeles, CA. These cities are unique insofar as they offer data that are sufficiently rich to enable an evaluation of not only of depolicing, but also the character of police withdrawal. That we are able to replicate these analyses across four contexts – where most other work drawing on similar data is limited to a single city – render our findings relatively broad.

Our first task is to evaluate whether depolicing occurred following the BLM protests. Using a regression discontinuity-in-time approach, we find a discontinuous and persistent drop in officer contact with civilians. This finding is durable and holds across all contexts. Leveraging 911 calls in two cities, we find that the change in stops is not driven by citizen requests for assistance. Second, we establish competing hypotheses about the pro- or anti-social nature of this decline. If depolicing is pro-social, we expect to see improvements in the quality of activity that does occur: more contraband per search, more arrests per stop, and diminished racial disparities. On the other hand, if declines in activity are anti-social we expect no change in overall quality, or even perhaps the opposite (Nix, Wolfe, and Campbell, 2018). On balance, we observe an improvement in arrest rates in all contexts; diminished Black/white stop disparities in all but one context; and no consistent improvement in hit rates. We conclude that the character of depolicing is mixed, and context dependent. With respect to crime, the critical test is violence, which is understood to be less sensitive to policing tactics themselves relative to against property/society crimes. We do not find consistent evidence that increased violent crime accompanied depolicing following the protests.

We make several contributions: First, we offer systematic and robust evidence that protests can compel widespread and durable changes in police behavior. Second, we are unable to consistently link depolicing with either pro- or anti-social policing behaviors. This suggests that the nature of depolicing can be either structured or unstructured. For example, consistent pro-social changes in Seattle suggest that a structured effort by leadership may have occurred, promoting coherent reforms. Fully investigating the circumstances that lead to structured (versus unstructured) declines in police service provision is an area for future research. Finally, we show that depolicing does not have clear consequences for violent crime.

In what follows, we provide an overview of the existing literature on the extent to which public officials respond to protests. We then review evidence around the conditions under which depolicing is likely to occur, and how such declines in service may impact crime. We develop three hypotheses, concerning depolicing overall, the quality of depolicing, and the consequences of depolicing for public safety. We then describe our case selection, data and analytic strategy, and review the results.

Background

Can anti-police protests prompt depolicing?

Little research has explicitly evaluated whether anti-police protests are successful in extracting behavioral changes (and higher quality outcomes) from law enforcement. Instead, scholars have focused on peripheral questions, like the impact of anti-police protests on officers' morale (Deuchar, Fallik, and Crichlow, 2019; Mercado, 2019; Nix, Wolfe, and Campbell, 2018; Oliver, 2017), and downstream impacts on crime (Tiwari, 2016; MacDonald, 2019; Lohman, 2021; Capellan, Lautenschlager, and Silva, 2020). Yet, there is reason to think that anti-police protests could prompt altered behavior from police. Elected officials are responsive to protests organized around racial justice – especially when protests occur in the lawmaker's district (Gillion, 2012). Indirectly, protests impact lawmaker decisions by durably shifting public opinion (Wasow, 2020; Enos, Kaufman, and Sands, 2019). Unelected officials may likewise be moved out of a desire to protect the legitimacy of their institution, which may be of particular concern here, where law enforcement are facing a crisis of legitimacy that predated the events of 2020 (Bell, 2016; Meares, 2015; Alon-Barkat and Gilad, 2016). That police may pro-socially respond to protester concerns is reflected in findings that fatal interactions between police and black civilians declined following the Ferguson uprising in 2014 (Skoy, 2021).

Only a few researchers have directly examined the phenomenon of depolicing, and the evidence that it occurs systematically in response to external pressure is mixed. Interviews with officers themselves indicate that they believe depolicing happens, and that individuals engage in this behavior for a variety of reasons (Nix, Wolfe, and Campbell, 2018; Oliver, 2017; Gau, Paoline III, and Paul, 2022; Foster, Rossler, and Scheer, 2023). Scholars call withdrawal from duty that might occur in response to anti-police protests *dissent shirking*, where officers change their behavior because they feel that they have been unfairly maligned by the public (Chanin and Sheats, 2018; Eckhouse, 2022). Dissent shirking, however, carries with it the implication of retaliation, where officers withdraw from duty because they disagree with critiques of their activities. Officers may also alter their behavior because they do not want to draw attention to themselves or risk becoming the focus of a civil inquiry. This kind of behavior is better characterized as avoidant than dissident (Nix, Wolfe, and Campbell, 2018). Officers may likewise police less because they are overwhelmed by the demands of the job, and public criticism may exacerbate feelings of burnout (Oliver, 2017). Scholars leverage strain theory to organize officers responses to an increasingly stressful work environment that may result from external criticism (Nix, Wolfe, and Campbell, 2018). From this perspective, depolicing is a coping mechanism officers leverage to reduce stress by avoiding putting themselves in situations where they might use force, that invite evaluation, or to alleviate psychological distress arising from sustained criticism (Agnew, 1992; Paoline III, 2004; Paoline III, 2003; Mac Donald, 2017).

Existing research thus suggests that officers may engage in depolicing for a variety of reasons, but that this behavior is most likely to occur in response to a strained or high-stress work environment. However, because it may occur in an unstructured, highly individual way, depolicing may not always be observable in the aggregate. For example, surveys of law enforcement both before and after the 2014 Ferguson uprising suggest that declines in discretionary service provision are limited in scope and duration (Marier and Fridell, 2020; Cheng and Long, 2022). Likewise, Chanin and Sheats (2018) find no change in police behavior in response to policy reforms imposed by the Department of Justice when misconduct violations are exposed, nor does Koslicki (2022) observe changes to use of force practices by the Minneapolis police department after the death of George Floyd. Yet, evaluations of agencies in Missouri post-Ferguson find that misdemeanor arrests declined across the state the year following the protests (Shjarback et al., 2017; Powell, 2022).

Whether anti-police protests can compel durable change in officer behavior is thus an open question. Qualitative evidence suggests that how individual officers respond to anti-police protests varies widely, but that declines in discretionary activity are most likely to follow from instances of extraordinary work-place strain. The volatile nature of the protests in many cities, ongoing criticism of law enforcement, and efforts by local officials to reform policing practices that followed suggest that the context of the 2020 BLM protests created a highly strained environment – exactly the circumstances that might give way to depolicing. For these reasons, we develop the following hypothesis:

Hypothesis 1: There will be a discontinuous decline in discretionary policing activities following the 2020 BLM protests.

Can depolicing be characterized as pro- or anti-social?

As noted above, reasons for curtailing discretionary policing activities are varied. Officers may be concerned about the legitimacy of their institution and depolicing may reflect accountability and responsiveness to community demands. For example, Mummolo (2018) finds that directives from agency leadership to document more fully the reason for conducting a Terry stop in New York City yielded an immediate increase in high-quality stops that produce evidence of criminal activity. This suggests that structured directives aiming to improve service provision can indeed yield pro-social policing outcomes. Moreover, discretionary policing practices can be altered in day-to-day activities.

In the absence of a clear, top-down directive from agency leadership it is not possible to ascertain a singular motive for declines in discretionary policing using the kind of administrative data required to evaluate whether depolicing is occurring in the first place. However, even without assessing underlying motivations, we may be able to characterize the substantive nature of declining police activity as either pro- or anti-social. Pro-social depolicing would manifest as increasing efficiency (for example, higher hit rates when stops do occur as in Mummolo (2018)), declining racial disparities in stops, or better service provision in marginalized communities (Nix, Wolfe, and Campbell, 2018; Shjarback et al., 2017; Rosenfeld and Wallman, 2019).

The character depolicing is likely to take following a protest is unclear, and may be context specific. On one hand, extant literature suggests that protests can function to hold city officials accountable by exerting political pressure. Mayors and city councils often have a fair amount of control over local law enforcement activities, particularly via budgets. The city councils in all four cities included in this analysis – in keeping with most other major U.S. cities – passed resolutions to address use-of-force by law enforcement in the days following the onset of the protests. It may be the case that any decline in discretionary police activity we observe following the protests reflects accountability to protester demands vis-a-vis elected officials. In this instance we may expect the quality of policing to improve overall.

It may also be the case that declining police stops that follow the protests are accompanied by an improvement in the quality of policing overall because of the nature of tactics over which law enforcement have discretion. For example, preemptive policing practices require a high volume of civilian interactions in order to identify contraband, and relies on questionable assumptions to determine when crime is likely to occur (Epp, Maynard-Moody, and Haider-Markel, 2014; Epp and Erhardt, 2021; Meares, 2015). Engaging less in these kinds of stops may therefore lead to an improvement in quality of policing simply because officers shift to relying on practices requiring a higher threshold of suspicion (e.g. officers may shift to relying more heavily on probable cause rather than consent to initiate contact with citizens). Both of these possibilities – accountability and shifts in the kind of stops officers engage in – lead to the following expectation around the quality of policing we may observe, post-BLM:

Hypothesis 2a: There will be a discontinuous improvement in the quality of policing overall following the 2020 BLM protests.

However, the accuracy of this hypothesis is highly contingent on local political context, so the null hypothesis – that there will be no change in the quality of policing overall – is also plausible. In the event that officers are simply policing less without changing the manner in which they police, we might expect to see no change in measures of quality. Even as there is some evidence that the city councils in all four cities included in the analysis attempted to address use-of-force practices following the protests, there is not much evidence that these efforts were more than symbolic (Walsh, Goodin-Smith, and Seidman, 2021; Kamb and Beekman, 2021-12-14). Thus, we may observe declines in service that we would characterize as anti-social, yielding no real improvement in terms of contraband hit rates and the like. This generates the following alternative hypothesis:

Hypothesis 2b: There will not be a discontinuous change in the quality of policing overall following the 2020 BLM protests.

In communities where over-policing is a concern, declines in police stops might be welcome. Such declines can still be thought of as anti-social because they are less likely to be driven by efforts to improve policing outcomes, and may lead to declines in public safety.

Does depolicing lead to increased crime?

Much of the existing literature on depolicing examines the impact of anti-police protests on crime, where the fear is that protests compel police to withdraw, and the belief is that proactive policing from which they withdraw is vital to deterring (especially violent) crime (Capellan, Lautenschlager, and Silva, 2020). This has been dubbed *The Ferguson Effect*, since this line of thinking gained traction in the wake of the 2014 Ferguson uprising. Researchers have failed to clearly link both anti-police protests and depolicing to meaningful changes in violent crime rates (Tiwari, 2016; MacDonald, 2019; Lohman, 2021; Capellan, Lautenschlager, and Silva, 2020; Rosenfeld and Wallman, 2019).[1] This leads to the following, final hypothesis:

Hypothesis 3: There will not be a discontinuous change in violent crime following the 2020 BLM protests.

Data and Design

Case selection

To select cities for inclusion in our analysis, we surveyed the open data websites of the top 20 most populous cities in the United States and collected all available incident-level data related to policing or crime.[2] Then, we identified the cities that had the following data available: 1) incident-level records of police activity, such as stops and/or officer initiated 911 calls; 2) incident-level records including metrics of quality of policing, such as recovery of contraband, and crucially, the race of civilian stopped; 3) incident-level records of crime that we could aggregate to the daily level (where previous work has relied on monthly counts of crime provided by the UCR), and 4) incident-level records up to at least one year prior to the onset of the protests.[3] We identified four cities that met these criteria: Seattle, WA, Philadelphia, PA, Los Angeles, CA, and Austin, TX (Table 1).[4]

No other city of which we are aware provides data detailed enough to evaluate our

[1]The one exception is Piza and Connealy (2022), who found that violent crime rose in and around the east precinct in Seattle, which police vacated during the course of the 2020 Black Lives Matter protests. However, total withdrawal of police service provision, including in response to civilian calls for service, is an extreme form of depolicing that does not characterize draw-downs in discretionary service sometimes observed following public scrutiny.

[2]We consulted city employees involved in managing the city's data where appropriate.

[3]This is to have sufficient data to conduct temporal placebo tests and assess if the 2020 BLM protests had an effect on our outcomes of interest larger than pre-treatment discontinuities

[4]See Table A1 for a full enumeration of police data available for the top 20 cities.

Table 1: Data Availability Across Top 20 Most Populated US Cities

City	State	Population Size	Crime Data	Call Data	Stop Data	Stop Race Data	Mayor Party	Evidence Of BLM Protest
New York City	NY	8804190	✗	✓	✓	✓	Democrat	✓
Los Angeles	CA	3898747	✓	✓	✓	✓	Democrat	✓
Chicago	IL	2746388	✓	✗	✗	✗	Democrat	✓
Houston	TX	2304580	✗	✗	✗	✗	Democrat	✓
Phoenix	AZ	1608139	✓	✓	✗	✗	Democrat	✓
Philadelphia	PA	1608139	✓	✗	✓	✓	Democrat	✓
San Antonio	TX	1434625	✗	✗	✗	✗	Independent (Progressive)	✓
San Diego	CA	1386932	✗	✓	✓	✓	Republican	✓
Dallas	TX	1304379	✓	✗	✗	✗	Democrat	✓
San Jose	CA	1013240	✗	✓	✗	✗	Democrat	✓
Austin	TX	961855	✓	✗	✓	✗	Democrat	✓
Jacksonville	FL	949611	✗	✗	✗	✗	Republican	✓
Fort Worth	TX	918915	✓	✗	✗	✗	Republican	✓
Columbus	OH	905748	✗	✗	✗	✗	Democrat	✓
Indianapolis	IN	897041	✓	✗	✗	✗	Democrat	✓
Charlotte	NC	874579	✓	✓	✓	✓	Democrat	✓
San Francisco	CA	873965	✓	✓	✗	✗	Democrat	✓
Seattle	WA	737015	✓	✓	✓	✓	Democrat	✓
Nashville	TN	715884	✓	✓	✗	✗	Democrat	✓
Denver	CO	715522	✓	✗	✓	✗	Non-Partisan (Democrat)	✓
D.C.	N/A	712816	✓	✗	✓	✓	Democrat	✓

Note: Shaded rows denote cities included in study. Population data from U.S. Census (2020). Catalogue of available data conducted May 2023.

hypotheses. For example, Denver, CO did not provide information on the race of civilian stopped; Dallas and Phoenix do not provide any information that would allow us to evaluate the quality of policing; New York City does not provide crime data; Washington D.C.'s stop data do not extend a year prior to the 2020 BLM protest; and data from Charlotte, NC is aggregated at the monthly level, precluding a daily regression discontinuity-in-time design that helps mitigate omitted variable bias. The four cities included in our analysis provide some regional coverage, as well as variation in the intensity of the protests and the responses of local city officials. Although there is no variation by city partisanship (all 4 cities were governed by a Democratic mayor), only three of 20 cities have Republican mayors.[5] Therefore, the cities included in our analyses are characteristic of the vast majority of major American cities. Additionally, all of the 20 largest cities in the U.S. had BLM protests. Since there is no cross-sectional variation in exposure to protests, we assess within-city variation in our outcomes of interest before and after the BLM protests to effectively understand the

[5] San Diego is headed by a Republican, and has all requisite data but crime. We evaluate San Diego as a robustness check. The full analysis is included in Appendix Section I, and referenced where appropriate.

consequences of the protests, looking for patterns across cities.

The protests in each of our four cities were characterized by clashes between the police and protesters, which likely created a strained work environment for officers.[6] Although there is no major city that did not have a protest during the summer of 2020, there is variation in the intensity of the protests, which may in turn impact the likelihood and character of depolicing we might observe. Seattle perhaps represents the most volatile protest environment under study. The protests were contextualized by a long, conflicted history between community activists and law enforcement, which came to a boiling point in May of 2020. The city adopted a contract with the Seattle Police Officers Guild that pushed the department out of compliance with a previous consent decree, and moved to end outside monitoring imposed by that same decree ("Timeline of Seattle Police Accountability" 2021). The protests lasted long into the summer, were characterized by police violence towards citizens, and famously, officers abandoned the East Precinct ("Timeline of Seattle Police Accountability" 2021). The protests in Los Angeles were similarly intense, leading Governor Gavin Newsom to declare a state of emergency, deploying the National Guard (Reyes-Valarde et al., 2020-05-31; Petrie, 2020-0604). These two cases are perhaps where we would most expect to see depolicing occur and persist. In contrast, while protests in Austin and Philadelphia were also contentious, they died out by the end of the first week of June (Fernandez and Mccullough, 2020; Gammage, 2020). We expect to observe at least a short term reduction in discretionary policing practices, but may be unlikely to persist over the long term. In sum, although there are no major cities in the U.S. where protests did not occur, the cities included represent variation in protest intensity, which may impact outcomes of interest.

In the event that we observe depolicing in the cities under study, we do not have the ability to determine the mechanisms that lead to declines in discretionary tactics engaged

[6]There is no way to speak about the protests that occurred in the cities under study without noting that violence occurred. In all four cities under study there is evidence that officers engaged in violence towards protesters. In all four cities there is evidence that citizens on the street at times committed property damage during the protests. Whether any property damage that occurred during the protests can be attributed to individuals who self-affiliated with the protests is unclear from available evidence. That should not be taken to mean that the BLM movement encouraged violence.

by officers. The literature suggests that top-down directives from public officials can lead to depolicing that is pro-social in nature. In the absence of evidence of a top down directive, such analysis requires individual-level attitudinal information for each officer in each city, which is outside the scope of the present analysis. While we do not have evidence that agency leadership issued such directives in any city, there is variation across cities in how public officials responded to the protesters, which may provide suggestive context for the character that depolicing might take. In Austin and Los Angeles, city officials responded quickly and resolutely in support of the protesters' demands. In Austin, less than three weeks after the protests erupted, the City Council approved cutting law enforcement's budget by a third and passed a suite of policies designed to increase transparency and accountability (Venkataramanan, 2020; Fernandez and Mccullough, 2020). In Los Angeles, the City Council moved to cut the LAPD's budget by $150 million dollars, reallocating a sizable portion to non-police responses to non-violent emergencies and poverty relief (Munoz, 2021-03-03).

In contrast, while Seattle's Mayor was at first supportive of the protesters, the City Council was divided, and in the wake of the protests have continuously voted to increase funding for the police department ("Timeline of Seattle Police Accountability" 2021). Given these differences in Los Angeles and Austin relative to Seattle, we might expect depolicing in the first two cases to appear pro-social and in the latter case to appear anti-social. The response of public officials in Philadelphia was more mixed. The city council put forward proposals for an oversight commission and new restraints on the kind of force tactics available to officers (McCrystal, 2020), but overall, the city did not appear particularly interested in pressuring the department to undertake radical change (Walsh, Goodin-Smith, and Seidman, 2021).[7] As previously noted, the motivations for depolicing likely vary at the officer level, and depolicing may appear pro-social for a variety of reasons having mostly to do with

[7]How public officials in a given city responded to the protests likely further varied by the partisanship of city leaders, rendering the omission of cities with Republican leadership a weakness for this analysis. We might anticipate that responses supportive of police diminished the workplace strain felt by officers, decreasing the likelihood of depolicing, and perhaps especially antagonistic or anti-social depolicing. We evaluate San Diego to address this concern. Indeed, city officials in San Diego responded to the protests by voting almost unanimously to *increase* the police budget by $27 million (Flores, 2020-06-10).

the nature of discretionary practices themselves. Thus, questions around the quality that depolicing might take are thorny, tangled and difficult to study. Even so, context around protest intensity and the responses of public officials across cities both highlights that there is important variation across cases, and can help us interpret whatever findings we may have.

Data

To assess if the 2020 BLM protests reduced discretionary policing (*Hypothesis 1*), we draw on the following data in each city: traffic stops in Austin (January 2019-December 2020);[8] pedestrian stops (July 2018-February 2023) and traffic stops (July 2018-February 2023)[9] in Los Angeles; pedestrian (January 2018-December 2022) and traffic stops (January 2018-December 2022) in Philadelphia;[10] and Terry stops (March 2015-February 2022) in Seattle. We aggregate these data to a day-level time series characterizing the daily number of *stops*. If *Hypothesis 1* is correct, we would expect *stops* to decrease post-protest.

For *Hypothesis 2*, we evaluate whether the 2020 BLM protests changed policing quality. First, we assess if the 2020 BLM protests increased policing efficiency and reduced the rate of fruitless police-citizen contact. In each city, we use the stop data to construct a daily time series of two efficiency measures. *Hit rates* are the proportion of daily stops that result in the recovery of contraband.[11] In Seattle, *hit rates* are measured differently in that they are the proportion of daily stops that resulted in an arrest, citation, offense report, or referral for prosecution as opposed to a field contact without action taken, implying no identification of criminal wrongdoing (i.e. a fruitless stop). *Arrest rates* are the proportion of stops resulting

[8] Source: `https://data.austintexas.gov/browse?q=traffic+stops&sortBy=relevance&tags=racial+profiling`

[9] Source: `https://data.lacity.org/Public-Safety/LAPD-RIPA-AB-953-STOP-Person-Detail-from-7-1-2018-/bwdf-y5fe`

[10] `https://opendataphilly.org/datasets/vehicle-pedestrian-investigations/`

[11] In Philadelphia, contraband is “firearms,” “other weapons,” “narcotics,” or “other contraband” (Source: `https://www.phila.gov/media/20211109145453/executive-order-2021-06.pdf`). In Austin, contraband is “narcotics,” “illegal weapons,” and “gambling equipment” (Source: `https://www.phila.gov/media/20211109145453/executive-order-2021-06.pdf`). In Los Angeles, contraband is “firearms,” “ammunition,” “weapons other than a firearm,” “drugs/narcotics,” “alcohol”, “money,” “drug paraphernalia,” “cell phones,” “electronic devices,” “other contraband or evidence” (Source: https://policingequity.org/images/pdfs-doc/COPS-Guidebook22.pdf).

in an arrest, suggesting the identified offense during a stop was arrest-worthy. Our final measure of quality is changes in racially disparate stop patterns. To assess this, we evaluate if the 2020 BLM protests reduced the stop *rate ratio* between Black and white citizens. The *rate ratio* is the Black stop rate $((BlackStops/BlackPopulation) \times 10,000)$ divided by the white stop rate $((WhiteStops/WhitePopulation) \times 10,000)$.[12] If *Hypothesis 2a* is true, then the 2020 BLM protests will have a positive effect on *hit rates* and *arrest rates*, and a negative effect on the *rate ratio*. Conversely, if *Hypothesis 2b* is true, then the 2020 BLM protests will have no effect on *hit rates*, *arrest rates*, or the *rate ratio*.[13]

The policing quality outcomes we evaluate are by no means comprehensive. For instance, lawsuits, police-involved killings, complaints against officers, and budgetary shifts are other potential measures of policing quality. However, the outcomes we analyze are optimal for several reasons. Unlike metrics that reflect day-to-day police-citizen contact, many other policing quality outcomes cannot be measured at the daily-level due to their rarity (e.g. lawsuits, police-involved killings, budgetary shifts). Therefore, although lawsuits and police killings are high-profile, our outcomes measure the most quotidian instantiations of policing quality, which strike at the core of police-citizen relations. Moreover, since other measures of policing quality cannot be measured at the daily-level, we cannot assess the immediate, discontinuous impact of the protests, which poses problems for identifying the causal effect of BLM protests unperturbed by long-term secular factors. Second, certain policing quality measures, like complaints, are more endogenous to protest activity than police behavior measures. The public may be emboldened to report complaints against officers in response to BLM protests even if police malfeasance remains constant. Other dimensions of policing quality are important, but they pose analytical challenges that preclude causal estimation, and that do not represent day-to-day experiences civilians have with police, rendering them

[12] Racial group population estimates for each city are from the 2010 Census.

[13] To evaluate quality of policing, we focus our efforts on changes observed in vehicular stops (and omit pedestrian stops). We do this for parsimony, since each RDiT estimate presented requires a number of robustness checks, generating a lengthy and cumbersome appendix. An evaluation of pedestrian stops yields similar findings, and are available from the authors upon request.

poor choices for our specific inquiry.

To test *Hypothesis 3*, we use incident-level crime data obtained from each city's data portal. We rely on Federal National Incident Based Reporting System (NIBRS) rules for classifying crimes, separating them into three categories: *society* (e.g. drug possession, prostitution), *property* (e.g. burglary, car theft), and *violent* or *against persons* (e.g. robbery, assault). The temporal domain for the Austin,[14] Los Angeles,[15] Philadelphia,[16] and Seattle crime datasets[17] are January 2003-February 2022, January 2010-February 2023, January 2006-December 2022, and January 2008-February 2022 respectively. We are particularly interested in *violent crime*. This is because identification of violent crimes are less sensitive to police effort, and more reflective of civilian reporting (Rosenfeld and Wallman, 2019). Therefore, if police reduce activity post-protests, identification of against person crimes should be less endogenous to police response. In Seattle, for example, 94 percent of violent crimes are assault offenses. Five percent are (non-consensual) sex offenses. The rest are consensual sex offenses, homicide offenses, and human trafficking. To evaluate the effect of the 2020 BLM protests on *violent crime*, we generate a daily time series of the count of violent crimes. If *Hypothesis 3* is correct, the 2020 BLM protests should have no effect on *violent crimes*, although we may observe declines in the other two categories.

The independent variable for each of the daily time series is a binary indicator equal to one on or after the start of the 2020 *BLM protests* in each city. The start date for the BLM protests for Austin, Los Angeles, Philadelphia and Seattle is May 29, 2020;[18] May 28,

[14] `https://data.austintexas.gov/Public-Safety/Crime-Reports/fdj4-gpfu`
[15] `https://data.lacity.org/browse?q=crime&sortBy=relevance&tags=crime+data`
[16] `https://data.phila.gov/visualizations/crime-incidents`
[17] `https://data.seattle.gov/Public-Safety/SPD-Crime-Data-2008-Present/tazs-3rd5`
[18] `https://www.kxan.com/news/local/austin/demonstrators-arrested-overnight-at-austin-police-headquarters/`

2020;[19] May 30, 2020;[20] and May 29, 2020.[21]

Our data are ideal to test our hypotheses. Consistent with prior research (Shjarback et al., 2017; Powell, 2022), an alternative approach might use county-level data from the FBI Uniform Crime Report, and assess the differential effect of exposure to BLM protests on various crime and policing outcomes for agencies within specific counties. There are a few reasons to prefer our approach. First, not all police agencies report their crime and policing data to the FBI, and if they do, they do not necessarily report data for each month of a given year (30% of agencies do not report a full year's worth of data).[22] Our approach uses incident-level data that is directly reported from the agency instead of aggregated through an external organization (e.g. the FBI), reducing the risk of missing data. Second, our use of incident-level, daily data, allows us to assess the immediate, discontinuous effect of the BLM protests, reducing the risk that secular trends or events (e.g. COVID policies) will bias our coefficient estimates.

Estimation Strategy

We use a regression discontinuity-in-time (RDiT) design to assess the discontinuous effect of the BLM protests. The core identifying assumption is that no other events are driving police behavior outside the *BLM protests* (i.e. the *continuity assumption*). Given that we use daily-level data and an estimation strategy that allows us to assess the effect of the *BLM protests* at the point at which they begin, it is unlikely other factors are jointly driving the onset of the protests and shifts in police tactics. Although the RDiT design only allows us to assess immediate effects at the moment the *BLM protest* occurs, we believe this is the optimal research design since immediate effects are less likely to be perturbed by long-term

[19] `https://www.cbs8.com/article/news/local/black-lives-matter-protesters-take-to-los-angeles-streets-freeway-over-death-of-george-floyd/509-56517320-da5f-48ee-848c-8953efaec162`

[20] `https://www.inquirer.com/news/philadelphia/live/george-floyd-protest-philadelphia-minneapolis-police-20200530.html`

[21] `https://www.capitolhillseattle.com/2020/05/seattle-defiant-walk-of-resistance-protest-planned-over-george-floyd-killing/`

[22] See: `https://ucrbook.com/county-level-ucr-data.html`

unobserved secular trends that may influence policing and crime.

Importantly, given the *BLM protests* occur during the COVID-19 pandemic, our daily-level data in tandem with the RDiT design circumvents the possibility governmental and public COVID responses (e.g. restrictions) explain our results. Stay-at-home orders were initially implemented on March 17, 19, 23, and 24 of 2020 for Philadelphia, Los Angeles, Austin, and Seattle respectively, roughly two months before the *BLM protest* onset. Since the RDiT evaluates the immediate, discontinuous effect of the BLM protest at the daily-level, our *BLM protest* coefficients are likely not perturbed by concomitant COVID responses. Although the COVID-19 pandemic was underway during the *BLM protests*, its influence should be constant given the nature of the design. We confirm this through a temporal placebo test (referenced appropriately throughout the manuscript).

One potential shortcoming of our design is that the 2020 *BLM protests* characterize a *bundled treatment*. Mass and police behavior shifted across a variety of dimensions immediately after the onset of the *BLM protests* (e.g. some people participate in protests, some people stay home, the police counter-mobilize). We do not view this as a weakness primarily for theoretical reasons. Protests are never inherently clean or isolated treatments. By design, the mass public and police will immediately respond simultaneously to protests in a variety of different ways. Moreover, tactical policing shifts in response to protests are fundamentally interrelated and do not occur in a vacuum (Epp and Erhardt, 2021). Therefore, evaluating the effect of a protest always requires acknowledging the existence of concomitant responses, especially at the moment the protest begins.

Nevertheless, we address the bundled treatment problem by assessing longer-term effects of the *BLM protest* that are less likely to be affected by the immediate influence of mass mobilization and counter-mobilization on part of the police. To this end, we interpret RDiT coefficients after removing outcome data 1-100 days immediately after the *BLM protest.* Although this analysis may be subject to secular temporal trends and should be understood as descriptive, we believe it is necessary to understand the durability of some of the effects

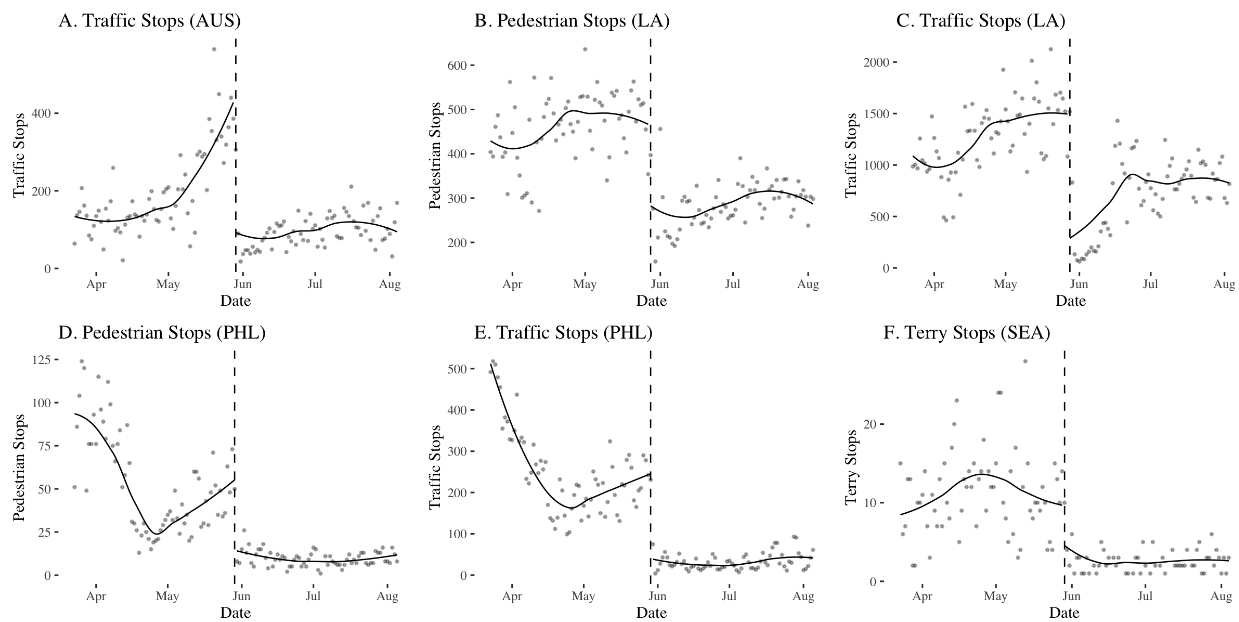

Figure 1: Policing Activity 2 Months Before and After BLM Protests. Each plot characterizes the amount (y-axis) of daily (x-axis) policing activity for Austin (Panel A), Los Angeles (Panel B-D), Philadelphia (Panels E-F), and Seattle (Panels G-H). Dashed vertical line denotes the onset of the 2020 BLM protests. Facet title denotes the specific outcome.

we observe. The absence of durable effect patterns may suggest our initial RDiT coefficients are highly idiosyncratic to the immediate consequences of the *BLM protest*.

For brevity, we interpret and present standardized RDiT coefficients using a uniform kernel, first-order polynomial (degree = 1), and mean-squared optimal bandwidth acquired with the `rdrobust` package in R (Calonico, Cattaneo, and Titiunik, 2015). We reference alternative specifications in the appendix as we describe the results when appropriate.

Results

Hypothesis 1: Depolicing

We find support for *Hypothesis 1.* Figure 1 describes the volume of discretionary policing activity before and after the protests. There is a clear, large, and immediate decrease across all measures of discretionary policing in every city under study.

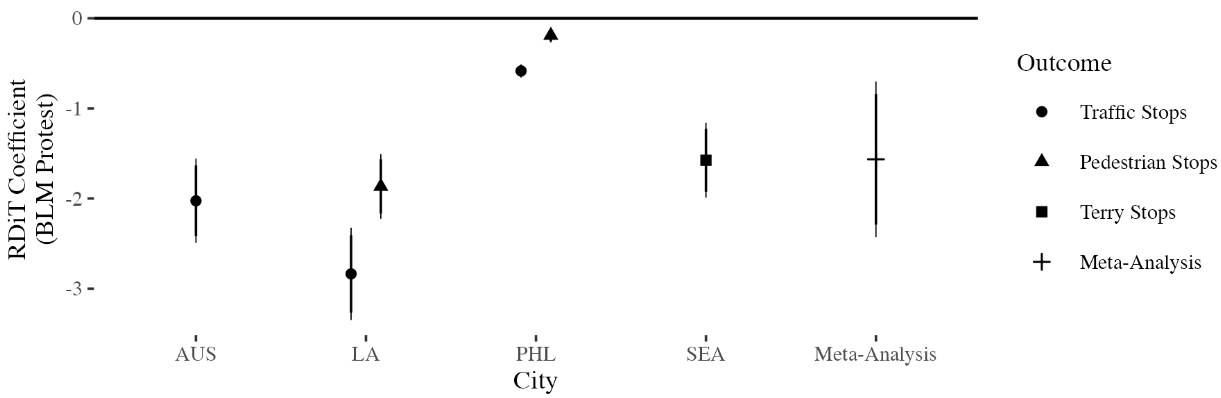

Figure 2: Standardized RDiT Coefficients Characterizing Effect of BLM Protests (y-axis) on Policing Activity Across Cities (x-axis). Shape denotes outcome type across the cities. All estimates are from RD specifications with a uniform kernel and polynomial degree equal to 1. Study-adjusted random effects meta-analytic coefficient on display. 95% CIs displayed derived from robust SEs. Associated regression estimates can be found in Appendix Table B2.

Figure 2 displays RDiT coefficients characterizing these relationships.[23] Across all cities and outcomes, there is a substantially large and statistically significant decrease in policing ($p < 0.001$ for all coefficients). The RDiT *BLM protest* coefficient ranges from -0.2 to -2.8 standard deviations. The study-adjusted random effects meta-analytic coefficient is -1.5 standard deviations.[24] These effects are not simply short-term effects intrinsic to the onset of the *BLM protests*. We also re-estimate RDiT coefficients omitting 1-100 days immediately after the protest to evaluate whether the discontinuous decrease in policing activity persists several days after it's initial onset. Observed decreases in police activity uniformly persist at least 100 days after the first BLM protests (Appendix Figures G36-G41). These estimates are also robust across kernel and polynomial specifications (Appendix Figure E2 - Figure

[23]Did a decline in police activity occur in ways that were similar across different neighborhoods? We may observe declining police activity in poorer or nonwhite neighborhoods, or we may observe shifting service provision from white and wealthy neighborhoods to nonwhite and poor neighborhoods. While we do not have geographic indicators associated with stops in all city contexts, we do have police beat where stops occurred in Seattle. We evaluated changes in Terry stops and officer initiated 911 calls among police beats with the highest/lowest concentrations of nonwhites, and among those beats where income fell above/below the city's median. We found no differences in depolicing by race, class, and geography in Seattle. The analyses are displayed in Table C6 in the appendix.

[24]"Study-adjusted" means if a city has more than one coefficient estimate due to having multiple outcomes in the data, the average of these coefficient estimates is taken within-city for the purposes of inclusion in the meta-analysis. This prevents "double counting" city estimates in the meta-analysis.

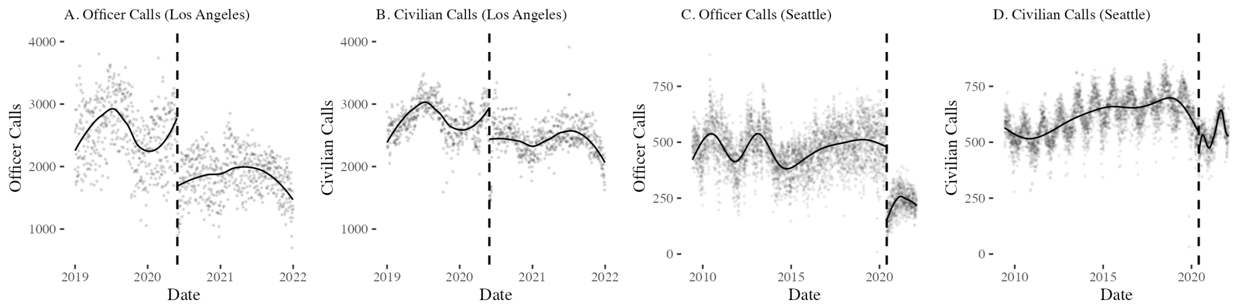

Figure 3: Officer and Civilian Emergency Calls (y-axis) Over Time (x-axis) in Seattle (Panels A-B) and Los Angeles (Panels C-D)

E5), and alternative bandwidths (Appendix Figure F6 - Figure F11). Finally, we conducted a temporal placebo test to assess whether changes in policing following the *BLM protests* were distinguishable from changes in policing behavior that may have occurred in all pre-*protest* days 30 days before the protest and 30 days after the beginning of the temporal domain of the data. Evidence of depolicing is robust to this test (Appendix Figures H66 - H69), therefore, secular events, such as the onset of COVID-19 and the respective lockdowns in each city, are not driving our findings.[25]

Is Depolicing Due To Reduced Civilian Demand?

An alternative explanation for the finding that the *BLM protests* decreased police activity is that civilians reduced demand for police services instead of the police restraining their activity. Reductions in civilian demand may be due to individuals staying home during the protest or a reticence to request police intervention brought on by the protests themselves (Ang et al., 2021). To assess this we leverage 911 call data from two of our four cities: Los Angeles (2019-01-01 to 2021-01-01) and Seattle (2010-01-01 to 2021-01-01).[26] Emergency call data from these cities can be disaggregated between calls initiated by civilians and by police

[25] As noted, we also evaluated depolicing in San Diego, CA, which features Republican leadership. Recall that the San Diego Mayor and City Council supported law enforcement, increasing their budget following the protests. Accordingly, while we do observe a decline in police activity directly following the protests, it returns to normal levels by the end of June. The full analysis is presented in Section I of the Appendix, and Appendix Figure I70.

[26] Source: `https://data.lacity.org/browse?q=calls%20for%20service&sortBy=relevance` and `https://data.seattle.gov/Public-Safety/Call-Data/33kz-ixgy`

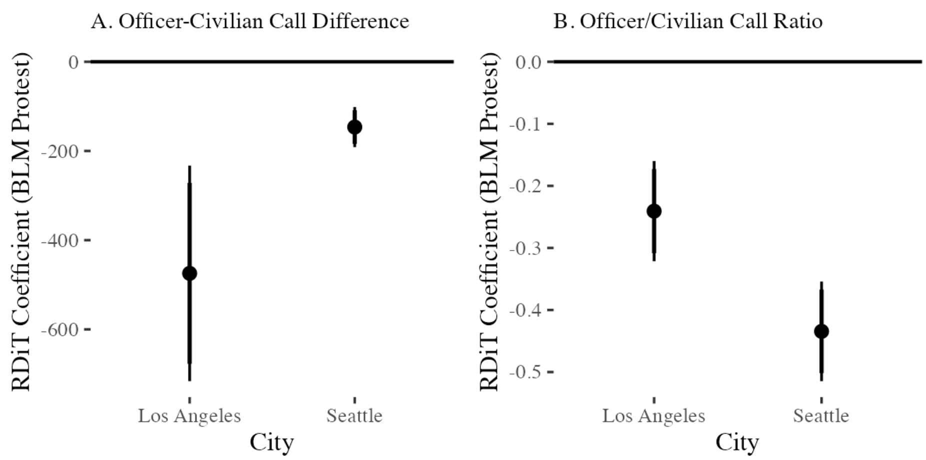

Figure 4: Assessing If Reductions in Policing Activity are Driven by Civilian Demand. The y-axis is the post-*BLM protest* RDiT coefficient, the x-axis is the city at use. All estimates are from RD specifications with a uniform kernel,polynomial degree equal to 1, and mean-squared optimal bandwidth selection (Calonico, Cattaneo, and Titiunik, 2015). 95% CIs displayed derived from robust SEs. Regression tables can be found in Appendix Table B3.

officers. Officer initiated 911 calls are often reports of encountered incidents.[27] Therefore, officer-initiated 911 calls serve as a measure of policing, while civilian-initiated calls serve as a measure of civilian demand for police services. Importantly, our goal is to rule out the possibility that declines in police stops are not wholly accounted for by reduced civilian demand. If we can show that decreases in officer-initiated 911 calls are more substantial and persistent than decreases in civilian-initiated calls, then we have evidence that police restraint is operative in policing patterns net of civilian demand.

Figure 3 shows officer and civilian calls over time. In Los Angeles, officer calls discontinuously and persistently decrease while civilian calls discontinuously decrease, but to a lesser extent then officer calls (Panels A-B). Likewise, in Seattle, officer calls appear to discontinuously decrease post-*BLM protest*, and the decrease persists well into 2020 (Panel C). Conversely, civilian calls discontinuously decrease only slightly post-*BLM protest* (Panel D), and rebound to the pre-protest mean by the end of 2020. Additionally, decreases in officer calls appear more substantial at the discontinuity than decreases in civilian calls.

[27] Based on our correspondence with LA and Seattle Open Data.

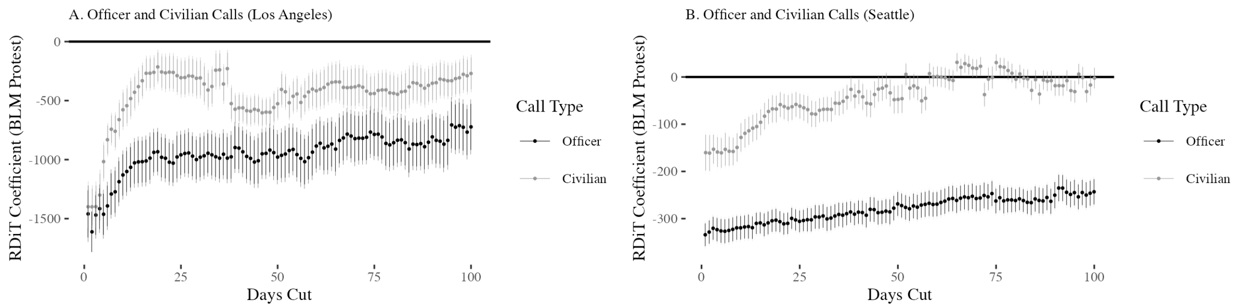

Figure 5: Assessing Persistence of Reductions in Civilian Demand and Police Activity. The x-axis is the number of days cut from right-hand side of the discontinuity in the data (but keeping days after intact). The y-axis is the post-*BLM protest* coefficient. Color denotes call type. Estimates are from RD specifications with a uniform kernel and polynomial degree equal to 1. 95% CIs displayed derived from robust SEs.

We conduct a formal test demonstrating that the discontinuous decrease in officer calls post-*BLM protest* was more substantial than the reduction in civilian calls. We assess the discontinuous decrease in the *difference* and *ratio* between officer and civilian calls (Figure 4). If the decrease is negative, then observed reductions in police activity are likely driven primarily by police themselves. We find statistically significant and substantial discontinuous reductions in the officer-civilian call *difference* (β = -0.4, SE = 0.04, $p < 0.001$; β = -0.24, SE = 0.04, $p < 0.001$) and officer/civilian call *ratio* (β = -146, SE = 22, $p < 0.001$; β = -434, SE = 123, $p < 0.001$) in both Los Angeles and Seattle.

Finally, we assess whether the decrease in officer calls post-*BLM protest* was more *persistent* than the reduction in civilian calls. To do this, we cut data from a specified number of days (1-100) immediately post-*BLM protest*, but keep all data after the cut number of days intact. In both cities, officer calls discontinuously and persistently decrease. But, civilian calls rebound in Los Angeles roughly 20 days post-*BLM protest* and do so within 50 days in Seattle (Figure 5). Coefficient difference tests suggest the decrease in officer calls is statistically lower than the decrease in civilian calls after cutting 1-100 days immediately post-*BLM protest* ($p < 0.05$ in all cases with the exception of the first four days cut in Los Angeles). In summary, consistent with *Hypothesis 1*, reductions in police activity are present even after accounting for citizen demand

Hypothesis 2: Pro- or Anti-Social Police Responses?

We find mixed evidence with respect to *quality of policing*. Recall that *Hypothesis 2a* anticipates an improvement in the quality of policing overall, while the null *Hypothesis 2b* anticipates no change (or a decline) in quality. We measure quality of policing in terms of change in hit rates, arrest rates, and Black/white stop rate ratios. Figure 6, Panel A suggests the *BLM protests* discontinuously increased the hit rate in Austin and Seattle by 0.04 and 0.15 respectively ($p < 0.001$, $p < 0.05$), 178% and 43% of the pre-treatment mean. However, the hit rate does not discontinuously increase post-*BLM protest* in Los Angeles or Philadelphia. Moreover, the positive coefficients for Austin and Seattle are not temporally sustained. Auxiliary analyses excluding days immediately post-*BLM protest* demonstrates improved hit rates last only 15 and 30 days for Austin and Seattle respectively post-*BLM protest* (Figures G42, G45). On balance, with respect to hit rates, we find support for the null hypothesis. Our interpretation of the results in Figure 6, Panel A, is consistent with the statistically insignificant random effects meta-analytic coefficient across the four cities (β = 0.021, SE = 0.013, $p = 0.12$).

In contrast, Panel B suggests the *BLM protests* discontinuously increased the arrest rate in every city. RDiT coefficients range from 0.03-0.15 ($p < 0.001$ for all cities except Seattle at $p < 0.05$), equivalent to 190-420% of the pre-treatment mean across the cities. The discontinuous improvement in arrest rates following the protests is robust to a variety of model kernel and polynomial specifications (Figures E2 - E5), and alternative bandwidths (Figures F16 - F19). Unlike the hit rate outcome, auxiliary analyses cutting days immediately post-*BLM protest* and re-estimating the RDiT coefficient suggests the improvement in arrest rates persists over time, even up to 100 days post-*BLM protest* (Figures G46-G49). These findings are informative, because they suggest the discontinuous increase in arrest rates is not simply a feature of police arresting more people participating in a protest conditional on initiating police contact. It is worth noting that in Austin and Los Angeles, we observe a dramatic improvement in arrest rates directly following the protest, which then declines

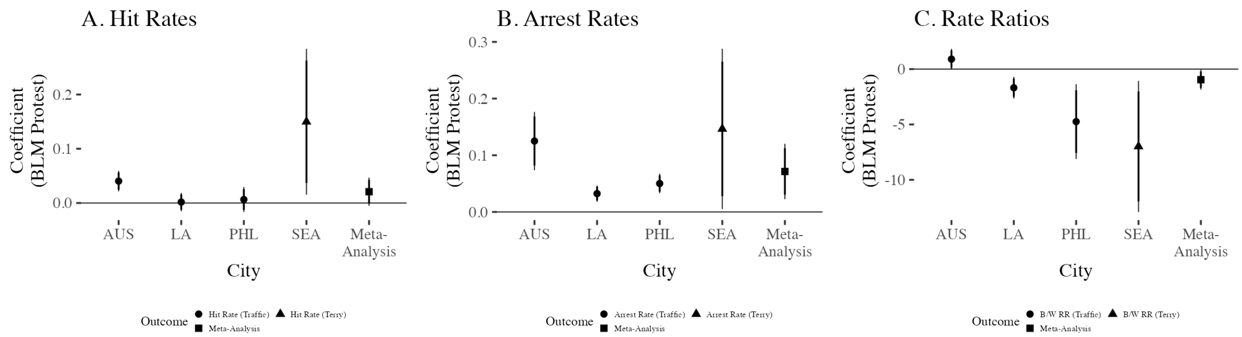

Figure 6: RDiT Estimates Characterizing Effect (y-axis) of BLM protests on Policing Quality Across Cities (x-axis). Panels A, B and C characterize the discontinuous effect of the BLM protests on *hit rates*, *arrest rates*, and *rate ratios* between whites and Black people. Shape denotes outcome type. All estimates are from RD specifications with a uniform kernel and polynomial degree equal to 1. Random effects meta-analytic coefficient on display for hit rate, arrest rate, and rate ratio outcomes. 95% CIs displayed derived from robust SEs. Associated regression estimates can be found in Appendix Table B4.

precipitously by 15 days after the onset of the protests, even as they remain statistically higher than prior to the protests over the longer term. The improvement in arrest rates across various specifications likewise passes the temporal placebo test (Figures H66 - H69). Durable and reliable improvements in arrest rates provide the strongest evidence that declines in police stops produced prosocial outcomes, supporting *Hypothesis 2a*. Our interpretation of the results on Figure 6, Panel B is consistent with our meta-analytic estimate suggesting the *BLM protest* discontinuously increased the arrest rate on average across the four cities ($\beta = 0.07$, SE = 0.02, $p < 0.01$).

Panel C indicates Black/white stop rate ratios discontinuously declined in Los Angeles, Philadelphia and Seattle post-*BLM protest*, with coefficients of -1.8 (Los Angeles, $p < 0.001$), -4.7 (Philadelphia, traffic, $p < 0.01$), and -7 (Seattle, Terry, $p < 0.05$), equivalent to 35%, 128%, and 98% of the pre-treatment mean respectively. However, Black/white stop *rate ratios* discontinuously increased post-*BLM protest* in Austin by 0.9 ($p < 0.05$). Auxiliary analyses cutting 0-100 days immediately post-*BLM protest* suggests the decrease in the Black/white stop rate ratio lasts at least up to 50 days (Figures G51 - G53). These estimates are most reliable across various specifications in Seattle and Philadelphia (Figures E4 - E5).

They are somewhat sensitive to model specification in Los Angeles (Figure E3), where it appears that the improvement is shorter term, occurs closer to the discontinuity, and returns to pre-treatment levels (Figure F22) sooner than in Seattle and Philadelphia. In contrast, in Austin the Black/white traffic stop rate ratio increased, though the increase lasted only 10 days, suggesting the discontinuous post-*BLM protest* coefficient is characterizing an effect that is short-term and intrinsic to the context of the protest (Figure G50). Evidence around quality of policing as measured by rate ratios is therefore mixed: declines in police stops coincided with an improvement in Black/white stop rate ratios in three out of four cities, and endured in two. In keeping with this interpretation, the meta-analytic estimate indicates the *BLM protest* discontinuously decreased the Black/white stop rate ratio by -0.96 (SE = 0.45, $p < 0.05$) – a statistically significant but substantively small coefficient.

Overall, we find the strongest evidence in support of *Hypothesis 2a* in Seattle and Philadelphia. Declining police stops did not produce durable improvements in hit rates in either city, but are associated with reliable improvements in both arrest rates and Black/white stop rate ratios that persist over time. In contrast, in Los Angeles and Austin, declining police stops were not accompanied by durable and reliable improvements in either hit rates or Black/white rate ratios, cannot be characterized as pro-social, and providing support for *Hypothesis 2b*.[28]

Hypothesis 3: Crime

Hypothesis 3 posits that there will be no change in violent crime following the protests. We also evaluate crimes against society and property, for comparison. The descriptive impact of the protests on crime is displayed in Figure 7. In each city it appears that violent crime dipped directly following the protests, but then resumed an overall upward trend that predated the unrest. Figure 8 displays the standardized RDiT coefficients characterizing the

[28]We also evaluated changes to the quality of policing in San Diego following the protests. It is unclear what to expect in terms of quality, conditional on partisanship of city leadership. We do observe findings similar to those observed in other cities: there is no impact of the protests on hit rates, but arrest rates and racial disparities do improve slightly. The full results are listed in Section I, Appendix Figure I70

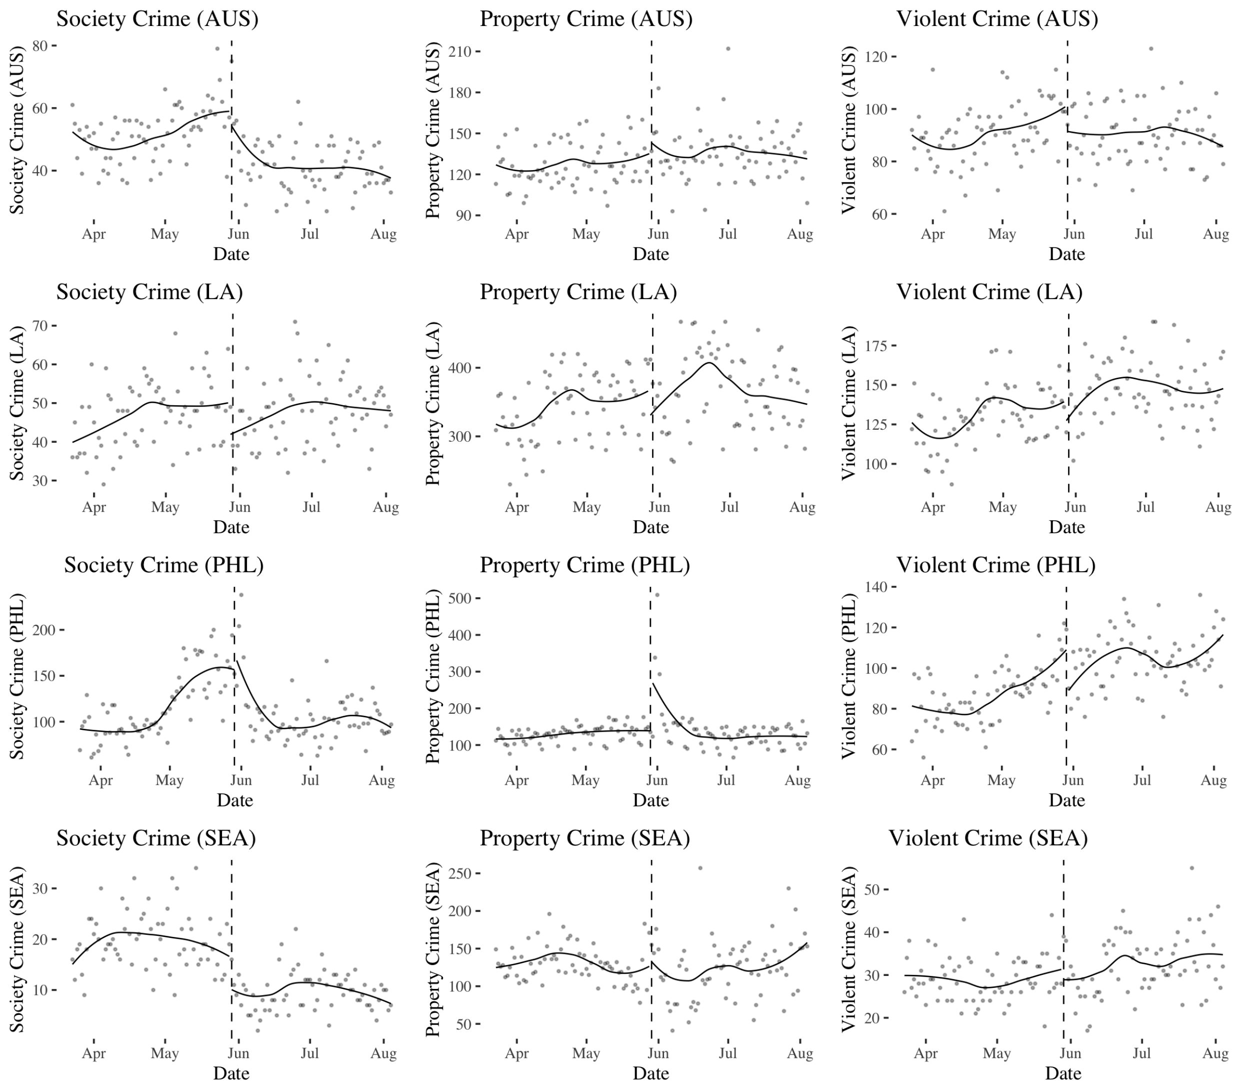

Figure 7: Crime 2 Months Before and After 2020 BLM Protests. The x-axis is the date, the y-axis is the crime type. For each row, the crime types are *society*, *property*, and *violent* from left to right. From top to bottom, each row characterizes data from Austin, Los Angeles, Philadelphia, and Seattle respectively. Dashed vertical line denotes the onset of the *BLM protests*. Loess models fit on each side of the *BLM protest* onset. Associated regression estimates can be found in Appendix Table B5.

discontinuous effect of the *BLM protest* on crime. Violent crime appears to increase in Philadelphia and LA by 0.5 ($p < 0.05$) and 0.9 ($p < 0.001$) respectively, but does not change in Austin and Seattle. In Austin and Seattle, the null effect of the protests on violent crime appears to be robust across model specifications (Figure E2 and Figure E5) and bandwidth specifications (Figure F24 and Figure F33), and is not distinguishable from patterns of violent

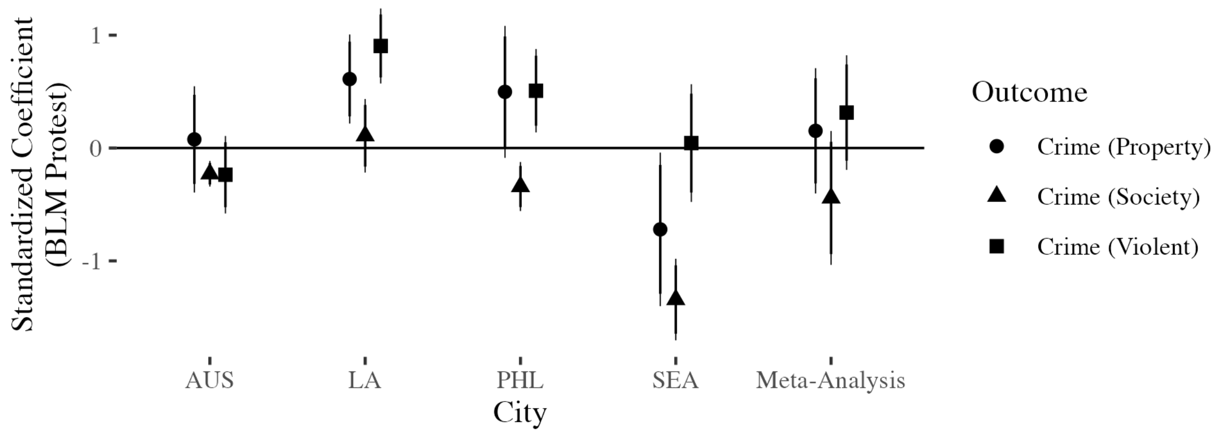

Figure 8: RDiT Estimates Characterizing Standardized Effect (y-axis) of *BLM Protests* on Crime Across Cities (x-axis). Shape denotes outcome type. All estimates are from RD specifications with a uniform kernel and polynomial degree equal to 1. Study-adjusted random effects meta-analytic coefficient on display. 95% CIs displayed derived from robust SEs.

crime occurring during the same time period the previous year (Figure H66 and Figure H69).

In Philadelphia and Los Angeles, the increases in violent crime following the protests appear to be a function of trends that pre-dated the protests. In both cities, the effect of the protests are not significant when the polynomial degree is quadratic or cubic (Figure E3 and Figure E4), suggesting that there is no change in violent crime close to the discontinuity (confirmed by an examination of alternative bandwidth specifications, Figure F27 and Figure F30). In Philadelphia, changes in violent crime reflected in the linear estimate are not distinguishable from the temporal placebo test, suggesting that factors other than the protests account for the upward trend (Figure H68). In Los Angeles, the difference between changes in violent crime that occurred around the protest and that which occurred the year prior approach statistical significance, but again do not hold across multiple polynomial degrees (Figure H67). We therefore cannot conclude that the protests themselves (and co-occurring declines in police activity) are responsible for increasing violent crime. Thus, across all city contexts we find support for Hypothesis 3.

We also evaluate changes in crimes against society and property, which existing literature suggests may fluctuate, given that they are more sensitive to actions taken by police themselves. The protests do not prompt change in crimes against society. Figure 8 displays the standardized RDiT coefficients characterizing the discontinuous effect of the *BLM protest* on crime. In all cities but Los Angeles, the linear RDiT coefficients suggest that crimes against society decrease overall. However, only in Seattle are shifts in this category of crime robust to various specifications (Figure E5), and distinguishable from fluctuations that occurred during the same time period the previous year (Figure H69). Across all contexts, changes to crimes against society are short term (Figures G54 - G65). On balance, we interpret the discontinuous effect of the *BLM protest* on crimes against society to be null.

Only in Philadelphia does it appear that the *BLM protest* led to a short term rise in property crime, descriptively. Figure 8 suggests that this temporary increase is not distinguishable from zero. Otherwise, property crime does not appear to change in Austin, increases in Los Angeles by 0.6 ($p < 0.001$), does not statistically change in Philadelphia, and decreases in Seattle by -0.72 ($p < 0.05$) after the BLM protest. In no city context is any observed change to property crime reliable across model specifications or persistent over time. We interpret the discontinuous effect of the *BLM protest* on property crime to be negligible.

In sum, we do not find robust and reliable evidence that the protests prompted a rise in any category of crime, including violent crime (the critical test). Meta-analytic estimates of the post-*BLM protest* effect on crime corroborate our interpretation of the results. On average, the meta-analytic, discontinuous effect of the *BLM protest* on property, society, and violent crimes is statistically insignificant (β = -0.44, SE = 0.30, p = 0.14; β = 0.15, SE = 0.28, p = 0.59; β = 0.31, SE = 0.26, p = 0.23). Contrast this with estimates concerning *Hypothesis 1*, which were highly robust, revealing, across all four contexts and multiple measures, a consistent and dramatic decline in police activity that is robust to a variety of robustness checks. We cannot be similarly confident in any of the findings around crime and are therefore unable to reject Hypothesis 3, which posits that the *BLM protests* will not discontinuously impact violent crime.[29]

Conclusion

We asked: What was the impact of the 2020 BLM protests on policing and public safety? In the event that the protests prompted declines in service provision, what quality does that depolicing take? And finally, did the protests and concurrent declines in police activity impact crime? In order to address these questions, we evaluate police activity in four cities, drawing together an array of data unprecedented in detail and breadth, and leverage an RDiT approach to identify the direct impact of the protests on downstream outcomes. Across all four cities, we find strong evidence that the 2020 BLM protests led to depolicing, but little evidence that declines in service provision were accompanied by a rise in violent crime.

With respect to the quality of policing, results are mixed. We do not observe any sustained improvement in hit rates. But, at the same time, we do observe an improvement in arrest rates, suggesting that when officers do stop people they are more often doing so for reasons related to observed criminal activity. Both declining stops and improved arrest rates are likewise accompanied by declining disparities in stop rates between Black and white civilians in three out of four cities, and improvements in racial disparities persist in two. We find stronger support for *Hypothesis 2a* in Seattle and Philadelphia, leading us to characterize the quality of depolicing in these cities as mostly pro-social. We find stronger support for *Hypothesis 2b* in Los Angeles and Austin, leading us to characterize the quality of depolicing in these cities as mostly anti-social. In all four cities, however, there was some evidence along one or more dimension that the character of depolicing was pro-social. More generally, less contact between police and civilians that does not impact public safety is normatively pro-social.

We cannot disentangle the mechanisms by which declines in service provision occur, and by extension the character depolicing takes. It may be that officers are genuinely improving

[29] We are not able to evaluate crime in San Diego, due to lack of appropriate data.

the deployment of stops in response to demands made by the protesters. There is not much contextual evidence to support this idea. The response from elected officials across cities was mixed, with the exception of Los Angeles where the Mayor and City Council were unified in support of the protester's demands. It may simply be that shifting to relying more heavily on practices that require a higher threshold of suspicion itself produces more pro-social outcomes rather than relying more heavily on tactics that have a lower threshold. This would comport with research elsewhere evaluating the impact of reliance on consent searches on downstream outcomes, which finds that these kinds of strategies do not improve public safety outcomes (Boehme, 2023; Epp and Erhardt, 2021). In sum, while there is evidence that depolicing yields some pro-social outcomes, contextual evidence and existing literature suggests this is because of the intrinsic nature of the stops themselves, and not a reflection of accountability to protester demands.

Our conclusions are three-fold. First, even though we cannot determine that officers reduced discretionary stops out of an interest in meeting protester demands, we nevertheless conclude that public protest is a viable path for citizens fighting to achieve a decrease in police-citizen interactions. In this regard, protesters were remarkably effective, causing a dramatic decline in police activities. This is an important finding as there has been much scrutiny of high-contact and high-discretion modes of policing that drive racial disparities but produce very little in terms of contraband, arrests, or other readily apparent crime-fighting benefits. That police made fewer stops across all four city contexts would likely be viewed as good news by the citizens calling for reforms in the summer of 2020.

Second, a chief contribution of our analysis concerns not only whether reduced contact occurred, but also how to characterize the nature of that reduced contact. We evaluated the quality of depolicing in terms of efficiency of stops, whether an arrest was made following a stop, and whether racial disparities improved. We therefore leverage new metrics of quality to develop a more nuanced understanding of withdrawal of service provision. Our findings suggest that this withdrawal can produce a net good, insofar as it is not associated with declining public safety. That said, identifying the city- or leadership-level factors that can promote systematic improvements in policing quality is an important area for future research.

Finally, our analysis offers reassurance to those worried about the public safety consequences of less policing. Violent crime, in particular, only appeared to increase in Los Angeles and Philadelphia, but these estimates do not stand up to rigorous analysis and appear to be attributable to secular trends not intrinsic to the protests themselves. In Seattle and Austin, violent crime did not change as a consequence of declining police activity. This finding highlights that the kind of discretionary police activities that can easily change in the day-to-day are not the kind of activities that most effectively reduce violent crime, giving cause to rethink rote policing practices in American cities.

References

Agnew, Robert (1992). "Foundation for a general strain theory of crime and delinquency". In: *Criminology* 30.1, pp. 47–88.

Alon-Barkat, Saar and Sharon Gilad (2016). "Political control or legitimacy deficit? Bureaucracies' symbolic responses to bottom-up public pressures". In: *Policy & Politics* 44.1, pp. 41–58.

Ang, Desmond et al. (2021). "Police violence reduces civilian cooperation and engagement with law enforcement". In: *HKS Working Paper No. RWP21-022*. URL: https://scholar.harvard.edu/ang/publications/police-violence-reduces-civilian-cooperation-and-engagement-law-enforcement.

Arango, Tim (2021). "A Year After George Floyd: Pressure to Add Police Amid Rising Crime". In: *The New York Times*. URL: https://www.nytimes.com/2021/05/23/us/george-floyd-police.html.

Bell, Monica C (2016). "Situational trust: How disadvantaged mothers reconceive legal cynicism". In: *Law & Society Review* 50.2, pp. 314–347.

Boehme, Hunter M (2023). "The influence of traffic stop policy changes on assaults against officers: A quasi-experimental approach". In: *Policing: A Journal of Policy and Practice* 17.1, paad002.

Calonico, Sebastian, Matias D Cattaneo, and Rocio Titiunik (2015). "rdrobust: An R Package for Robust Nonparametric Inference in Regression-Discontinuity Designs." In: *R J.* 7.1, p. 38.

Capellan, Joel A, Rachel Lautenschlager, and Jason R Silva (2020). "Deconstructing the Ferguson effect: A multilevel mediation analysis of public scrutiny, de-policing, and crime". In: *Journal of crime and justice* 43.2, pp. 125–144.

Chanin, Joshua and Brittany Sheats (2018). "Depolicing as dissent shirking: Examining the effects of pattern or practice misconduct reform on police behavior". In: *Criminal Justice Review* 43.2, pp. 105–126.

Cheng, Cheng and Wei Long (2022). "The effect of highly publicized police killings on policing: Evidence from large US cities". In: *Journal of Public Economics* 206, p. 104557.

Deuchar, Ross, Seth Wyatt Fallik, and Vaughn J Crichlow (2019). "Despondent officer narratives and the 'post-Ferguson'effect: exploring law enforcement perspectives and strategies in a southern American state". In: *Policing and Society* 29.9, pp. 1042–1057.

Eckhouse, Laurel (2022). "Metrics Management and Bureaucratic Accountability: Evidence from Policing". In: *American Journal of Political Science* 66.2, pp. 385–401.

Enos, Ryan D, Aaron R Kaufman, and Melissa L Sands (2019). "Can violent protest change local policy support? Evidence from the aftermath of the 1992 Los Angeles riot". In: *American Political Science Review* 113.4, pp. 1012–1028.

Epp, Charles R, Steven Maynard-Moody, and Donald P Haider-Markel (2014). *Pulled over: How police stops define race and citizenship*. University of Chicago Press.

Epp, Derek A and Macey Erhardt (2021). "The use and effectiveness of investigative police stops". In: *Politics, Groups, and Identities* 9.5, pp. 1016–1029.

Fernandez, Stacy and Jolie Mccullough (June 18, 2020). "Texas cities took quick actions after George Floyd's death. Advocates doubt they'll have a big impact". In: *Texas Tribune*. URL: https://www.texastribune.org/2020/06/18/texas-police-george-floyd/.

Flores, Jessica (2020-06-10). "San Diego defies calls to 'defund the police,' but it 'won't be business as usual' in California's second-largest city". In: *USA Today*. URL: https://www.usatoday.com/story/news/nation/2020/06/09/san-diego-police-budget-defund-calls/5330547002/.

Foster, Jacob T, Michael T Rossler, and Charles Scheer (2023). "An empirical analysis of depolicing behavior". In: *Police Practice and Research*, pp. 1–18.

Gammage, Jeff (June 8, 2020). "Peaceful, smaller protests take place across the region on Sunday; Rizzo mural is gone from the Italian Market." In: *Philadelphia Inquirer*. URL: https://www.inquirer.com/news/philadelphia-protests-george-floyd-rizzo-mural-italian-market-police-brutality-20200607.html#:~:text=News-

,Peaceful%2C%20smaller%20protests%20take%20place%20across%20the%20region%20on%20Sunday,for%20a%20ninth%20straight%20day..

Gau, Jacinta M, Eugene A Paoline III, and Nicholas D Paul (2022). "De-policing as a result of public scrutiny: examining officers' perceptions of negative media attention and willingness to engage in self-initiated activity". In: *Journal of Crime and Justice* 45.5, pp. 539–551.

Gause, LaGina (2022). "Revealing issue salience via costly protest: How legislative behavior following protest advantages low-resource groups". In: *British Journal of Political Science* 52.1, pp. 259–279.

Gillion, Daniel Q (2012). "Protest and congressional behavior: assessing racial and ethnic minority protests in the district". In: *The Journal of Politics* 74.4, pp. 950–962.

Kamb, Lewis and Daniel Beekman (2021-12-14). "After fatal shooting in protest zone, Seattle mayor's email called situation 'foreseeable and avoidable'". In: *Seattle Times*. URL: https://www.seattletimes.com/seattle-news/after-deadly-shooting-in-protest-zone-seattle-mayor-jenny-durkan-called-situation-foreseeable-and-avoidable-email-shows/.

Koslicki, Wendy M (2022). "Criticism does not constrain: testing for evidence of de-policing following the murder of George Floyd". In: *Policing: An International Journal* ahead-of-print.

Lohman, Robin Christine (2021). "Alterations in Crime and the Ferguson Effect: An Analysis of Crime Trends in the St. Louis, MO Metropolitan Area". PhD thesis. Walden University.

Mac Donald, Heather (2017). *The war on cops: How the new attack on law and order makes everyone less safe*. New York: Encounter Books.

MacDonald, John M (2019). "De-Policing as a Consequence of the So-Called Ferguson Effect". In: *Criminology & Pub. Pol'y* 18, p. 47.

Marier, Christopher J and Lorie A Fridell (2020). "Demonstrations, demoralization, and de-policing". In: *Criminology & Public Policy* 19.3, pp. 693–719.

McCrystal, Laura (June 12, 2020). "Philly City Council wants to ban police kneeholds and require new cops to live in the city." In: *Philadelphia Inquirer*. URL: https://www.inquirer.com/news/philadelphia-city-council-police-reform-bills-20200611.html.

Meares, Tracey (2015). "Broken windows, neighborhoods, and the legitimacy of law enforcement or why I fell in and out of love with Zimbardo". In: *Journal of Research in Crime and Delinquency* 52.4, pp. 609–625.

Mercado, Christopher (2019). "The Ferguson Effect in Contemporary Policing: Assessing Police Officer Willingness to Engage the Public". PhD thesis. City University of New York.

Mummolo, Jonathan (2018). "Modern police tactics, police-citizen interactions, and the prospects for reform". In: *The Journal of Politics* 80.1, pp. 1–15.

Munoz, Anabel (2021-03-03). "LA City Council allocates funds cut from LAPD to policing alternatives, homelessness prevention". In: *ABC7 Eyewitness News*. URL: https://abc7.com/defund-the-police-lapd-black-lives-matter-los-angeles-city-council/10384293/.

Nix, Justin, Scott E Wolfe, and Bradley A Campbell (2018). "Command-level police officers' perceptions of the "war on cops" and de-policing". In: *Justice quarterly* 35.1, pp. 33–54.

Oliver, Willard M (2017). "Depolicing: Rhetoric or reality?" In: *Criminal Justice Policy Review* 28.5, pp. 437–461.

Pagones, Stephanie (2020). "Crime surged amid George Floyd protests in parts of US". In: *Fox Business*. URL: https://www.foxbusiness.com/lifestyle/crime-protests-cities-george-floyd.

Paoline III, Eugene A (2003). "Taking stock: Toward a richer understanding of police culture". In: *Journal of criminal justice* 31.3, pp. 199–214.

— (2004). "Shedding light on police culture: An examination of officers' occupational attitudes". In: *Police quarterly* 7.2, pp. 205–236.

Petrie, Faith (2020-0604). "A Timeline of the George Floyd Protests in L.A." In: *The Los Angeles Sentinel*. URL: https://lasentinel.net/a-timeline-of-the-george-floyd-protests-in-la.html.

Piza, Eric L and Nathan T Connealy (2022). "The effect of the Seattle Police-Free CHOP zone on crime: A microsynthetic control evaluation". In: *Criminology & Public Policy* 21.1, pp. 35–58.

Powell, Zachary A (Nov. 2022). "De-policing, police stops, and crime". In: *Policing: A Journal of Policy and Practice* 17, paac070. URL: https://doi.org/10.1093/police/paac070.

Reyes-Valarde, Alejandra et al. (2020-05-31). "Looting hits Long Beach, Santa Monica as countywide curfew goes into effect". In: *The Los Angeles Times*. URL: https://www.latimes.com/california/story/2020-05-31/looting-vandalism-leaves-downtown-l-a-stunned.

Rosenfeld, Richard and Joel Wallman (2019). "Did de-policing cause the increase in homicide rates?" In: *Criminology & Public Policy* 18.1, pp. 51–75.

Shjarback, John A et al. (2017). "De-policing and crime in the wake of Ferguson: Racialized changes in the quantity and quality of policing among Missouri police departments". In: *Journal of criminal justice* 50, pp. 42–52.

Skoy, Evelyn (2021). "Black Lives Matter protests, fatal police interactions, and crime". In: *Contemporary Economic Policy* 39.2, pp. 280–291.

Smith, Savannah, Jiachuan Wu, and Joe Murphy (2020). "Map: George Floyd protests around the world". In: *CBS News*. URL: https://www.nbcnews.com/news/world/map-george-floyd-protests-countries-worldwide-n1228391.

"Timeline of Seattle Police Accountability" (2021). In: *American Civil Liberties Union*. URL: https://www.aclu-wa.org/pages/timeline-seattle-police-accountability.

Tiwari, Alisa (2016). "Protests, Policing, and Crime: An Analysis of Evidence Regarding the Ferguson Effect". In: *Journal of Public & International Affairs* 1.

Venkataramanan, Meena (Aug. 13, 2020). “Austin City Council cuts police department budget by one-third, mainly through reorganizing some duties out from law enforcement oversight”. In: URL: https://www.texastribune.org/2020/08/13/austin-city-council-cut-police-budget-defund/.

Walsh, Sean Collins, Oona Goodin-Smith, and Andrew Seidman (Apr. 22, 2021). “A guide to how policing laws in Philly and Pa. changed — or didn’t — after George Floyd’s murder”. In: *Philadelphia Inquirer*. URL: https://www.inquirer.com/news/philadelphia/derek-chauvin-george-floyd-philly-pa-police-laws-20210422.html#:~:text=Philadelphia%20City%20Council%20fast%2Dtracked,have%20more%20power%20and%20resources..

Wasow, Omar (2020). “Agenda seeding: How 1960s black protests moved elites, public opinion and voting”. In: *American Political Science Review* 114.3, pp. 638–659.

Supplemental Information

A Data Availability Across 20 Largest Cities

Table A1: Data Availability Across Top 20 Most Populated US Cities

City	State	Population Size	Crime Data	Call Data	Stop Data	Stop Race Data	Car Accident Data	Arrest Data	Shot Spotter Data	Use of Force Data	Complaint Data	Car Citation Data	SQF Data	Mayor Party	Evidence Of BLM Protest
New York City	NY	8804190	✗	✓	✓	✓	✓	✓	✗	✓	✓	✗	✓	Democrat	✓
Los Angeles	CA	3898747	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	Democrat	✓
Chicago	IL	2746388	✓	✗	✓	✓	✓	✓	✗	✗	✗	✗	✗	Democrat	✓
Houston	TX	2304580	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	Democrat	✓
Phoenix	AZ	1608139	✓	✓	✗	✗	✗	✓	✗	✓	✗	✓	✗	Democrat	✓
Philadelphia	PA	1608139	✓	✗	✓	✓	✗	✓	✗	✗	✗	✗	✗	Democrat	✓
San Antonio	TX	1434625	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	Independent (Progressive)	✓
San Diego	CA	1386932	✗	✓	✗	✓	✗	✗	✗	✗	✗	✗	✗	Republican	✓
Dallas	TX	1304379	✓	✗	✗	✗	✓	✓	✗	✗	✗	✗	✗	Democrat	✓
San Jose	CA	1013240	✗	✓	✗	✗	✓	✗	✗	✗	✗	✗	✗	Democrat	✓
Austin	TX	961855	✓	✗	✓	✗	✗	✓	✗	✗	✗	✗	✗	Democrat	✓
Jacksonville	FL	949611	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	Republican	✓
Fort Worth	TX	918915	✓	✗	✗	✗	✓	✗	✗	✗	✗	✗	✗	Republican	✓
Columbus	OH	905748	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	Democrat	✓
Indianapolis	IN	897041	✓	✗	✗	✗	✗	✗	✗	✓	✓	✗	✗	Democrat	✓
Charlotte	NC	874579	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	Democrat	✓
San Francisco	CA	873965	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	Democrat	✓
Seattle	WA	737015	✓	✓	✓	✓	✗	✓	✗	✗	✗	✗	✗	Democrat	✓
Nashville	TN	715884	✓	✓	✗	✗	✓	✗	✗	✗	✗	✗	✗	Democrat	✓
Denver	CO	715522	✓	✗	✓	✗	✓	✗	✗	✗	✗	✗	✗	Non-Partisan (Democrat)	✓
D.C.	N/A	712816	✓	✗	✓	✓	✓	✓	✓	✗	✗	✗	✗	Democrat	✓

Note: Shaded rows denote cities included in study. Population data from U.S. Census (2020).

B RDiT Tables Associated with Figures Presented in Text

Table B2: RDiT Coefficients Characterizing the Effect of BLM Protests on Policing Activities.

City	Stops	Coeff	SE	P-Val	N-Val	Effective N	Bandwidth Est.
Austin	Traffic	-2.03	0.24	0.000	514	56.41	28.21
LA	Pedestrian	-1.87	0.18	0.000	697	164.25	82.13
LA	Traffic	-2.84	0.26	0.000	697	117.37	58.68
Philly	Pedestrian	-0.19	0.04	0.000	880	124.09	62.04
Philly	Traffic	-0.58	0.04	0.000	880	91.81	45.91
Seattle	Terry	-1.58	0.21	0.000	1885	170.99	85.49

All estimates are specified with a uniform kernel and polynomial degree equal to 1. Standard errors are robust.

Table B3: RDiT Coefficients Characterizing Changes in Officer Civilian Initiated Calls Following the BLM Protest.

City	Officer:Civilian Calls	Coeff	SE	P-Val	N-Val	Effective N	Bandwidth Est.
LA	Call Difference	-146.43	22.96	0.000	4014	96.15	48.07
Seattle	Call Difference	-0.43	0.04	0.000	4014	118.86	59.43
LA	Call Ratio	-474.43	123.30	0.000	513	170.24	85.12
Seattle	Call Ratio	-0.24	0.04	0.000	513	227.64	113.82

All estimates are specified with a uniform kernel and polynomial degree equal to 1. Standard errors are robust.

Table B4: RDiT Coefficients Characterizing the Effect of BLM Protests on Policing Quality.

City	Outcome	Coeff	SE	P-Val	N-Val	Effective N	Bandwidth Est.
Austin	Hit rate	0.04	0.01	0.00	514	82.98	41.49
LA	Hit rate	0.00	0.01	0.88	697	409.87	204.94
Philly	Hit rate	0.01	0.01	0.85	2341	239.23	119.62
Seattle	Hit rate	0.15	0.07	0.02	1885	337.01	168.50
Austin	Arrest rate	0.13	0.03	0.00	514	62.42	31.21
LA	Arrest rate	0.03	0.01	0.00	697	339.91	169.95
Philly	Arrest rate	0.01	0.01	0.85	2341	239.23	119.62
Seattle	Arrest rate	0.15	0.07	0.02	1885	337.01	168.50
Austin	B/W rate ratio	0.90	0.48	0.04	357	64.22	32.11
LA	B/W rate ratio	-1.69	0.50	0.00	697	236.81	118.40
Philly	B/W rate ratio	-4.74	1.72	0.00	2337	209.27	104.64
Seattle	B/W rate ratio	-6.99	3.02	0.01	339	183.84	91.92

All estimates are specified with a uniform kernel and polynomial degree equal to 1. Standard errors are robust.

Table B5: RDiT Coefficients Characterizing the Effect of BLM Protests on Crime.

City	Crime Type	Coeff	SE	P-Val	N-Val	Effective N	Bandwidth Est.
Austin	Violent	-0.24	0.18	0.07	6358	276.61	138.31
Austin	Property	0.08	0.24	0.96	6358	138.35	69.18
Austin	Society	-0.23	0.06	0.00	6358	277.56	138.78
LA	Violent	0.90	0.17	0.00	3800	383.81	191.91
LA	Property	0.61	0.20	0.00	3800	189.13	94.56
LA	Society	0.11	0.17	0.26	3800	351.43	175.71
Philly	Violent	0.51	0.19	0.03	5263	172.81	86.40
Philly	Property	0.50	0.30	0.05	5263	284.67	142.33
Philly	Society	-0.34	0.11	0.00	5263	178.23	89.11
Seattle	Violent	0.04	0.27	0.80	4532	150.25	75.13
Seattle	Property	-0.72	0.35	0.02	4532	179.89	89.94
Seattle	Society	-1.34	0.18	0.00	4532	274.78	137.39

All estimates are specified with a uniform kernel and polynomial degree equal to 1. Standard errors are robust.

C Depolicing does not vary by demographic composition of beat

Table C6: Regression discontinuity coefficient difference

	Coefficient difference	P-value	Bandwidth	Measure	DV
(1)	−0.878	.950	25	Income	Terry stops
(2)	−0.974	.946	50	Income	Terry stops
(3)	−1.209	.934	100	Income	Terry stops
(4)	1.460	.806	25	Nonwhite	Terry stops
(5)	1.069	.860	50	Nonwhite	Terry stops
(6)	1.166	.852	100	Nonwhite	Terry stops
(7)	−16.107	.985	25	Income	Calls
(8)	−19.631	.979	50	Income	Calls
(9)	−22.031	.976	100	Income	Calls
(10)	13.948	.976	25	Nonwhite	Calls
(11)	17.272	.964	50	Nonwhite	Calls
(12)	20.131	.947	100	Nonwhite	Calls

D Efficiency Is Not a Function of Identifying More Criminal Activity

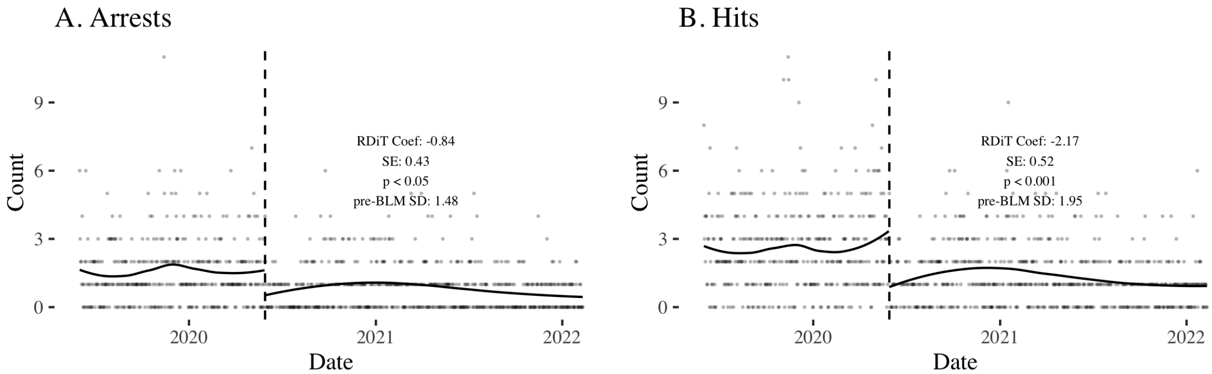

Figure D1: Terry Stop Arrest (Panel A) and Hit Counts (Panel B, y-axis) Over Time (x-axis). Loess lines fit on each side of the *BLM protest* discontinuity. Dashed vertical line denotes *BLM protest* onset. Annotations denote RDiT coefficients using a running variable to the 1st degree.

E Alternative RDiT Specifications

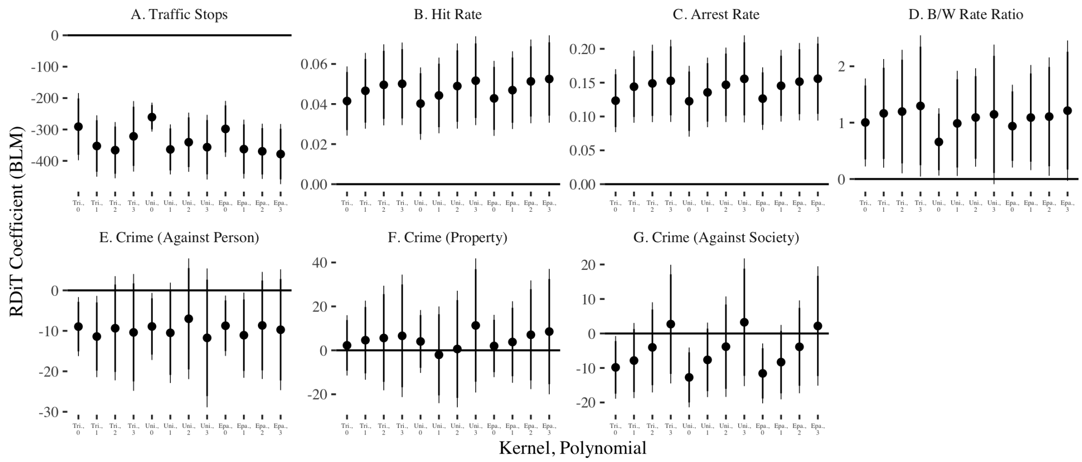

Figure E2: Alternative RDiT Specifications Across Outcomes (Austin) The x-axis characterizes the different kernel/polynomial specifications (0 = difference-in-means, 1 = linear polynomial, 2 = quadratic polynomial, 3 = cubic polynomial; Tri. = triangular kernel, Uni. = uniform kernel, Epa. = epanechnikov kernel). The y-axis characterizes the unstandardized RDiT coefficient for each of the respective outcomes (characterized by separate facets). 95% CIs displayed from robust SEs

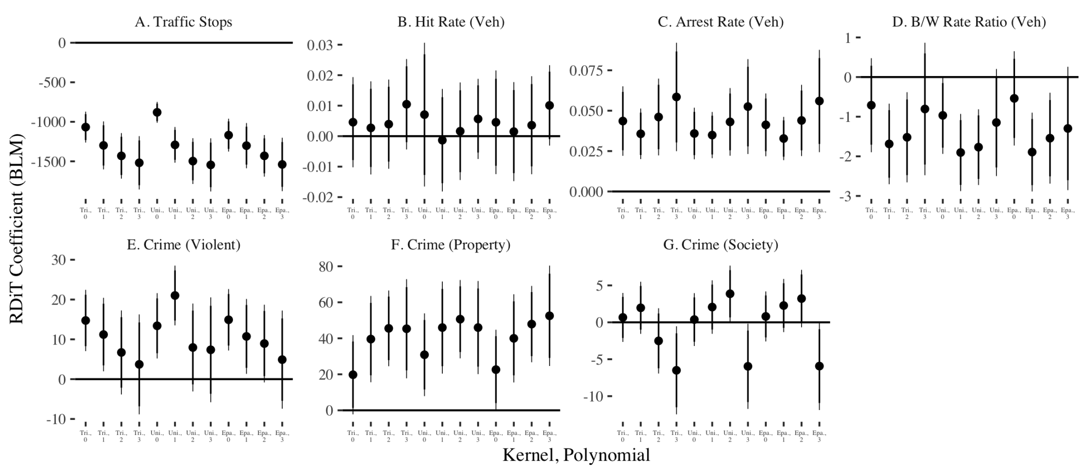

Figure E3: Alternative RDiT Specifications Across Outcomes (Los Angeles) The x-axis characterizes the different kernel/polynomial specifications (0 = difference-in-means, 1 = linear polynomial, 2 = quadratic polynomial, 3 = cubic polynomial; Tri. = triangular kernel, Uni. = uniform kernel, Epa. = epanechnikov kernel). The y-axis characterizes the unstandardized RDiT coefficient for each of the respective outcomes (characterized by separate facets). 95% CIs displayed from robust SEs

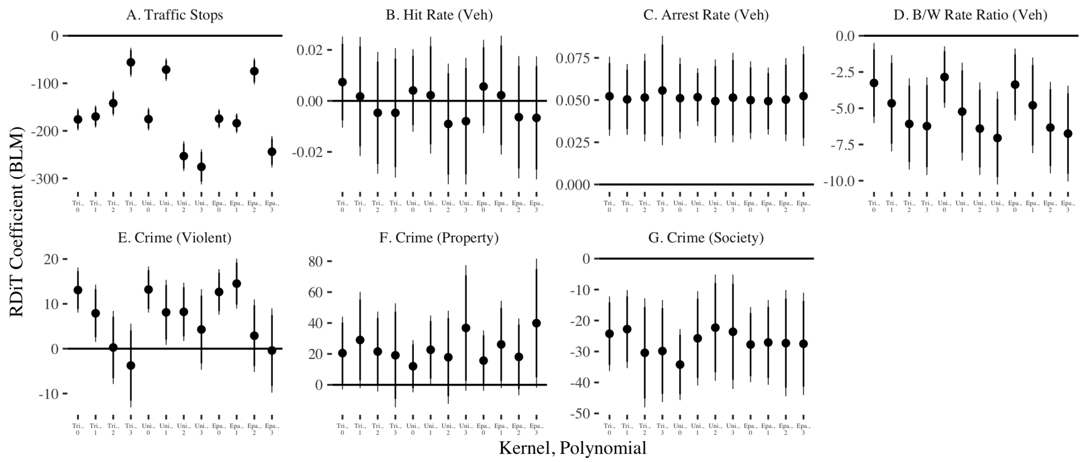

Figure E4: Alternative RDiT Specifications Across Outcomes (Philadelphia) The x-axis characterizes the different kernel/polynomial specifications (0 = difference-in-means, 1 = linear polynomial, 2 = quadratic polynomial, 3 = cubic polynomial; Tri. = triangular kernel, Uni. = uniform kernel, Epa. = epanechnikov kernel). The y-axis characterizes the unstandardized RDiT coefficient for each of the respective outcomes (characterized by separate facets). 95% CIs displayed from robust SEs

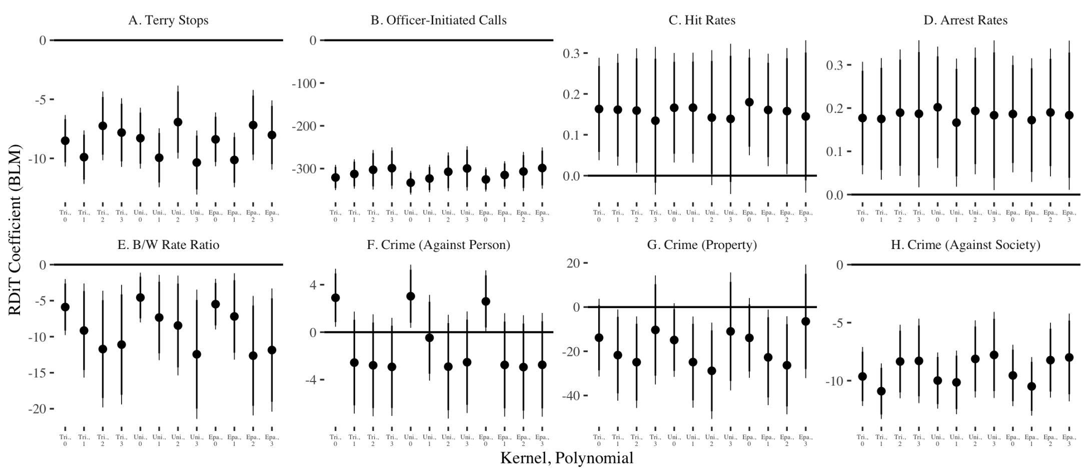

Figure E5: Alternative RDiT Specifications Across Outcomes (Seattle) The x-axis characterizes the different kernel/polynomial specifications (0 = difference-in-means, 1 = linear polynomial, 2 = quadratic polynomial, 3 = cubic polynomial; Tri. = triangular kernel, Uni. = uniform kernel, Epa. = epanechnikov kernel). The y-axis characterizes the unstandardized RDiT coefficient for each of the respective outcomes (characterized by separate facets). 95% CIs displayed from robust SEs

F Alternative Bandwidths

F.1 Policing Activity

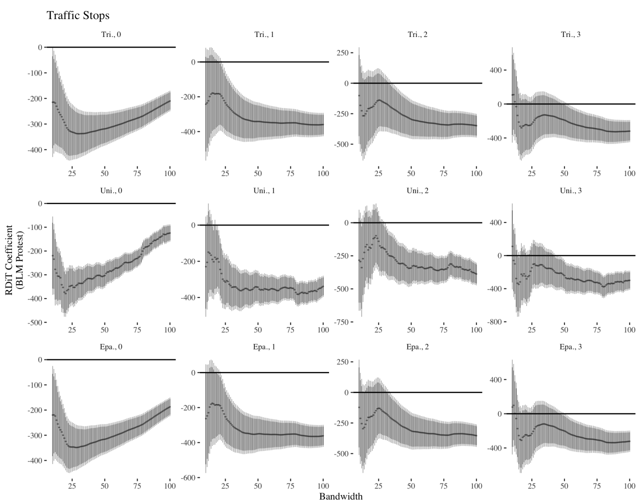

Figure F6: Alternative Bandwidths (10-100 Days, Austin, Traffic Stop Outcome) The x-axis is the number of days used in the data before and after the *BLM protest*. The y-axis is the unstandardized RDiT coefficient. Each facet characterizes a different kernel (Tri. = triangular, Uni. = uniform, Epa. = Epanechnikov) and polynomial (0, 1, 2, 3) specification. 95% CIs displayed from robust SEs.

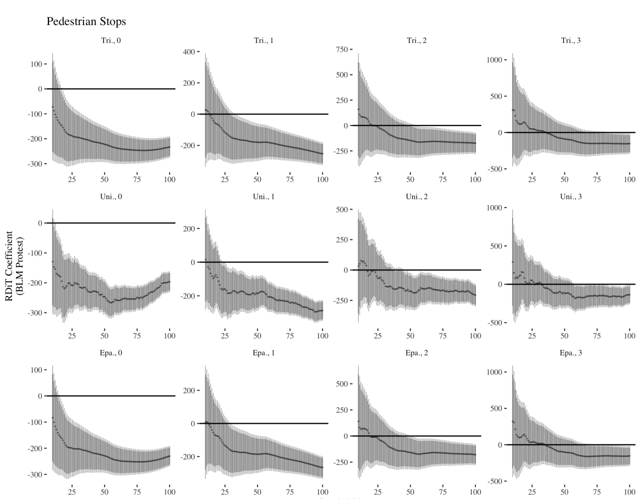

Figure F7: Alternative Bandwidths (10-100 Days, Los Angeles, Pedestrian Stops Outcome) The x-axis is the number of days used in the data before and after the *BLM protest*. The y-axis is the unstandardized RDiT coefficient. Each facet characterizes a different kernel (Tri. = triangular, Uni. = uniform, Epa. = Epanechnikov) and polynomial (0, 1, 2, 3) specification. 95% CIs displayed from robust SEs.

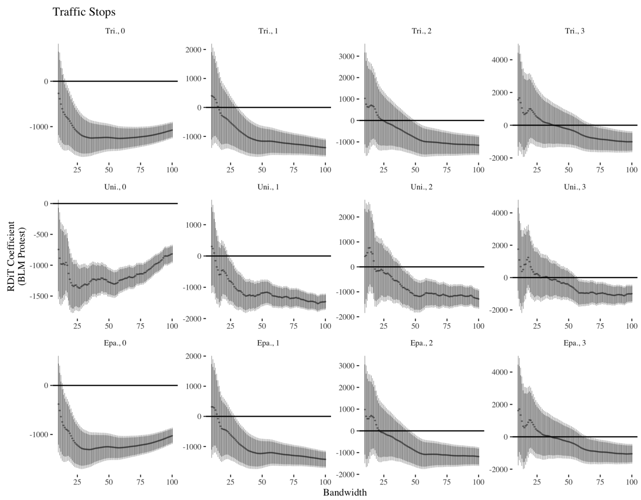

Figure F8: Alternative Bandwidths (10-100 Days, Los Angeles, Traffic Stops Outcome) The x-axis is the number of days used in the data before and after the *BLM protest*. The y-axis is the unstandardized RDiT coefficient. Each facet characterizes a different kernel (Tri. = triangular, Uni. = uniform, Epa. = Epanechnikov) and polynomial (0, 1, 2, 3) specification. 95% CIs displayed from robust SEs.

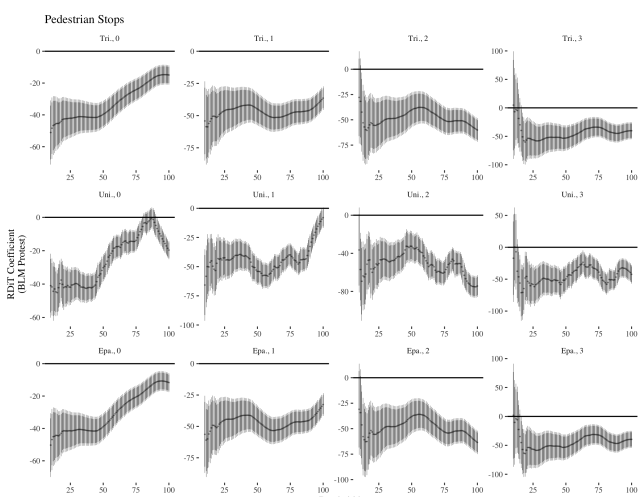

Figure F9: Alternative Bandwidths (10-100 Days, Philadelphia, Pedestrian Stops Outcome) The x-axis is the number of days used in the data before and after the *BLM protest*. The y-axis is the unstandardized RDiT coefficient. Each facet characterizes a different kernel (Tri. = triangular, Uni. = uniform, Epa. = Epanechnikov) and polynomial (0, 1, 2, 3) specification. 95% CIs displayed from robust SEs.

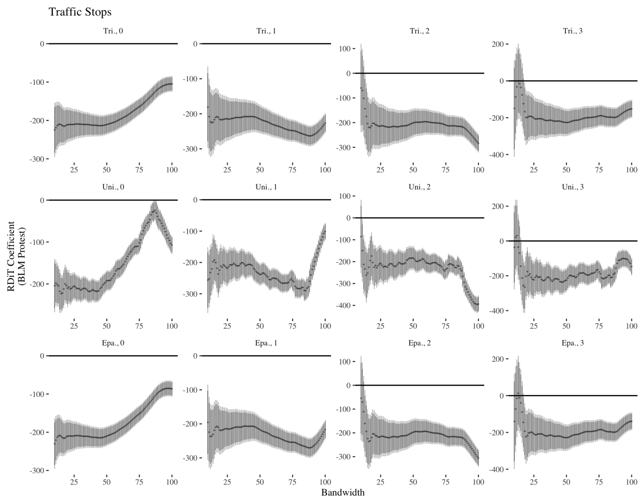

Figure F10: Alternative Bandwidths (10-100 Days, Philadelphia, Vehicle Stops Outcome) The x-axis is the number of days used in the data before and after the *BLM protest*. The y-axis is the unstandardized RDiT coefficient. Each facet characterizes a different kernel (Tri. = triangular, Uni. = uniform, Epa. = Epanechnikov) and polynomial (0, 1, 2, 3) specification. 95% CIs displayed from robust SEs.

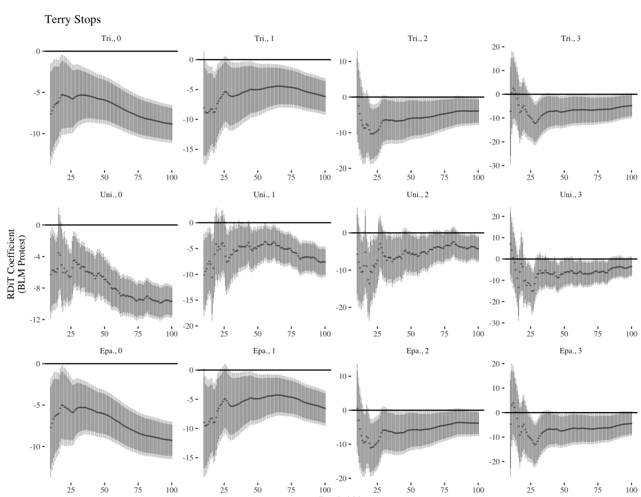

Figure F11: Alternative Bandwidths (10-100 Days, Seattle, Terry Stops Outcome) The x-axis is the number of days used in the data before and after the *BLM protest*. The y-axis is the unstandardized RDiT coefficient. Each facet characterizes a different kernel (Tri. = triangular, Uni. = uniform, Epa. = Epanechnikov) and polynomial (0, 1, 2, 3) specification. 95% CIs displayed from robust SEs.

F.2 Hit Rates

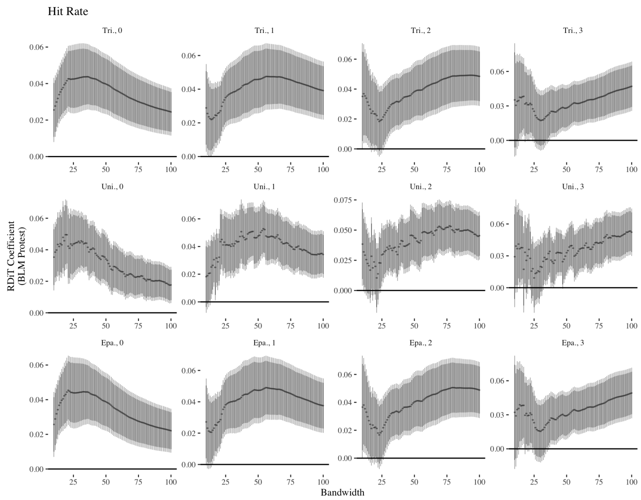

Figure F12: Alternative Bandwidths (10-100 Days, Austin, Hit Rates Outcome) The x-axis is the number of days used in the data before and after the *BLM protest*. The y-axis is the unstandardized RDiT coefficient. Each facet characterizes a different kernel (Tri. = triangular, Uni. = uniform, Epa. = Epanechnikov) and polynomial (0, 1, 2, 3) specification. 95% CIs displayed from robust SEs.

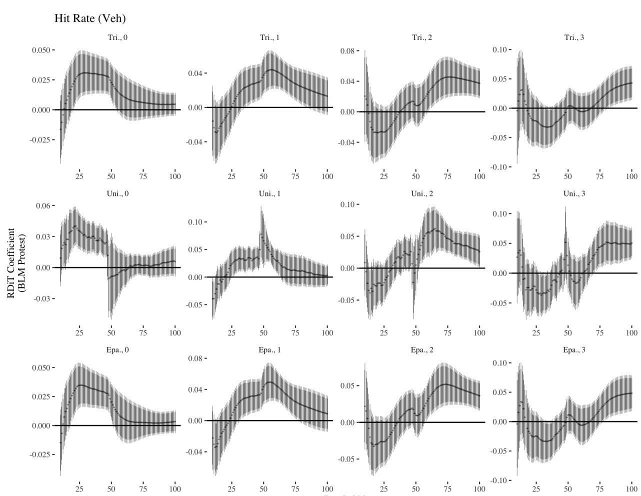

Figure F13: Alternative Bandwidths (10-100 Days, Los Angeles, Vehicle Hit Rates Outcome) The x-axis is the number of days used in the data before and after the *BLM protest*. The y-axis is the unstandardized RDiT coefficient. Each facet characterizes a different kernel (Tri. = triangular, Uni. = uniform, Epa. = Epanechnikov) and polynomial (0, 1, 2, 3) specification. 95% CIs displayed from robust SEs.

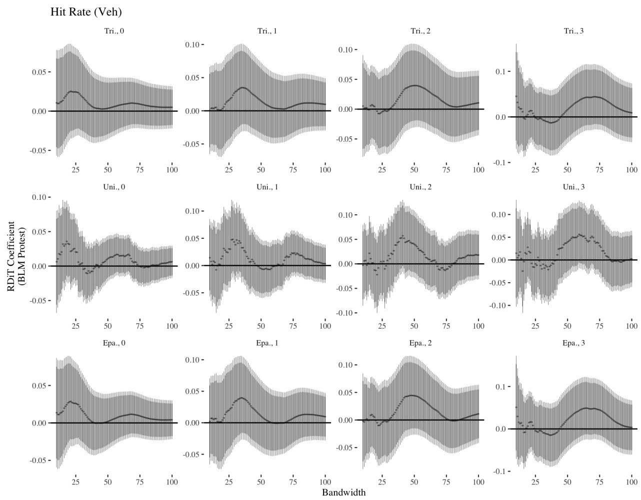

Figure F14: Alternative Bandwidths (10-100 Days, Philadelphia, Vehicle Hit Rates Outcome) The x-axis is the number of days used in the data before and after the *BLM protest*. The y-axis is the unstandardized RDiT coefficient. Each facet characterizes a different kernel (Tri. = triangular, Uni. = uniform, Epa. = Epanechnikov) and polynomial (0, 1, 2, 3) specification. 95% CIs displayed from robust SEs.

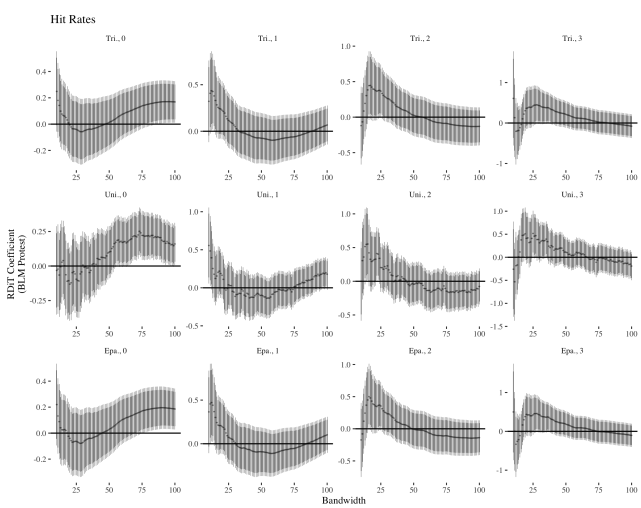

Figure F15: Alternative Bandwidths (10-100 Days, Seattle, Terry Stop Hit Rates Outcome) The x-axis is the number of days used in the data before and after the *BLM protest*. The y-axis is the unstandardized RDiT coefficient. Each facet characterizes a different kernel (Tri. = triangular, Uni. = uniform, Epa. = Epanechnikov) and polynomial (0, 1, 2, 3) specification. 95% CIs displayed from robust SEs.

F.3 Arrest Rates

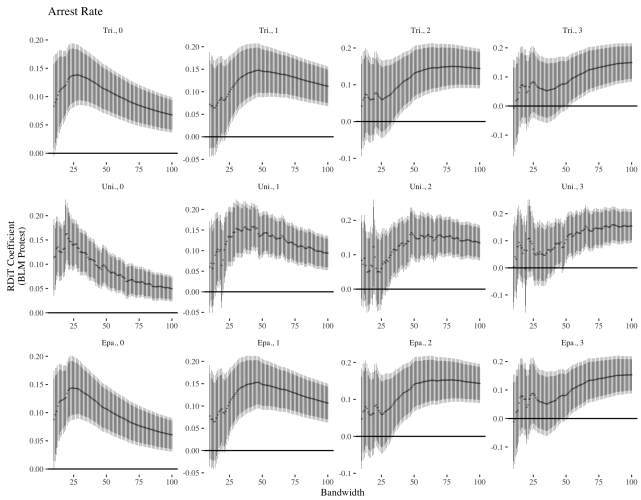

Figure F16: Alternative Bandwidths (10-100 Days, Austin, Vehicle Arrest Rates Outcome) The x-axis is the number of days used in the data before and after the *BLM protest*. The y-axis is the unstandardized RDiT coefficient. Each facet characterizes a different kernel (Tri. = triangular, Uni. = uniform, Epa. = Epanechnikov) and polynomial (0, 1, 2, 3) specification. 95% CIs displayed from robust SEs.

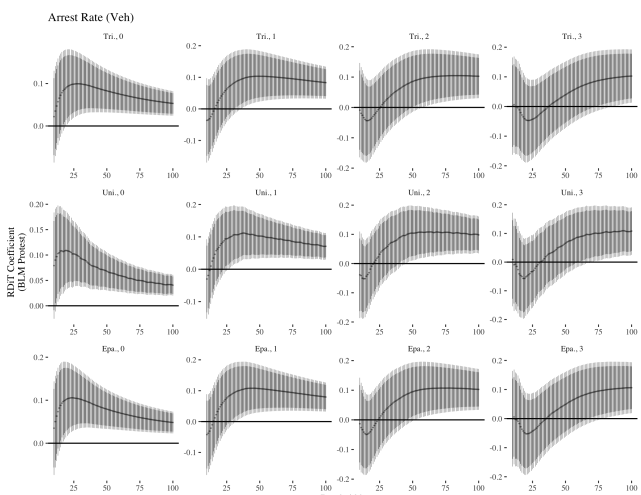

Figure F17: Alternative Bandwidths (10-100 Days, Los Angeles, Vehicle Arrest Rate Outcome) The x-axis is the number of days used in the data before and after the *BLM protest*. The y-axis is the unstandardized RDiT coefficient. Each facet characterizes a different kernel (Tri. = triangular, Uni. = uniform, Epa. = Epanechnikov) and polynomial (0, 1, 2, 3) specification. 95% CIs displayed from robust SEs.

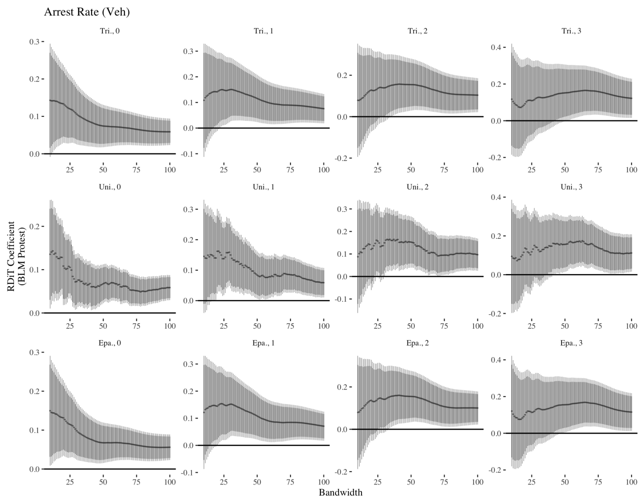

Figure F18: Alternative Bandwidths (10-100 Days, Philadelphia, Vehicle Arrest Hit Rates Outcome) The x-axis is the number of days used in the data before and after the *BLM protest*. The y-axis is the unstandardized RDiT coefficient. Each facet characterizes a different kernel (Tri. = triangular, Uni. = uniform, Epa. = Epanechnikov) and polynomial (0, 1, 2, 3) specification. 95% CIs displayed from robust SEs.

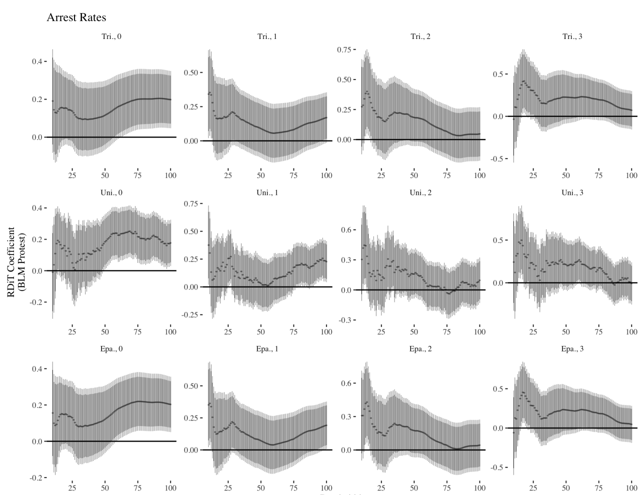

Figure F19: Alternative Bandwidths (10-100 Days, Seattle, Terry Stop Arrest Hit Rates Outcome) The x-axis is the number of days used in the data before and after the *BLM protest*. The y-axis is the unstandardized RDiT coefficient. Each facet characterizes a different kernel (Tri. = triangular, Uni. = uniform, Epa. = Epanechnikov) and polynomial (0, 1, 2, 3) specification. 95% CIs displayed from robust SEs.

F.4 Rate Ratios

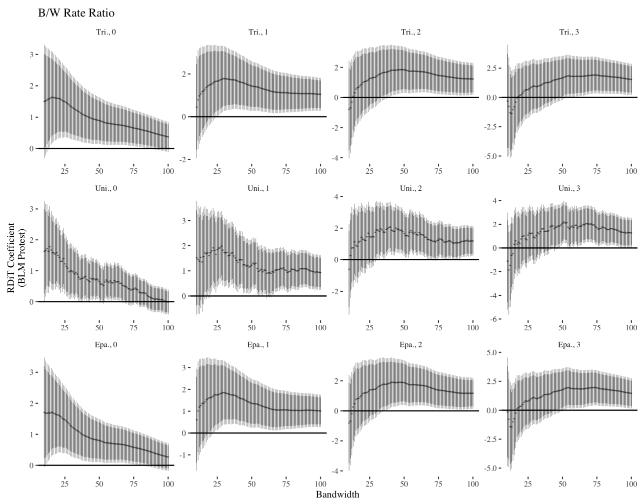

Figure F20: Alternative Bandwidths (10-100 Days, Austin, Black/white Vehicle Stop Rate Ratios Outcome) The x-axis is the number of days used in the data before and after the *BLM protest*. The y-axis is the unstandardized RDiT coefficient. Each facet characterizes a different kernel (Tri. = triangular, Uni. = uniform, Epa. = Epanechnikov) and polynomial (0, 1, 2, 3) specification. 95% CIs displayed from robust SEs.

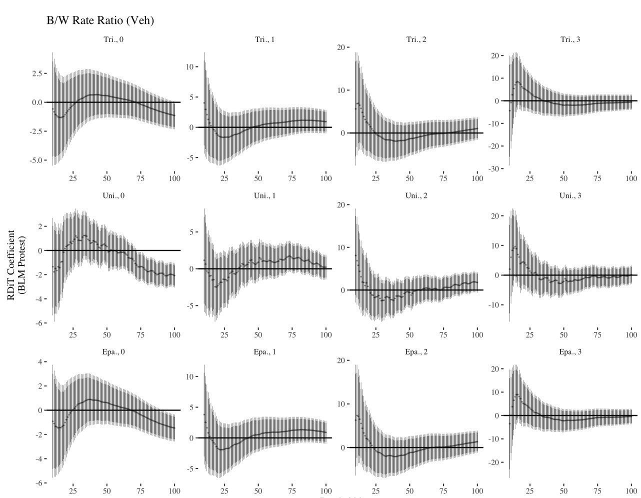

Figure F21: Alternative Bandwidths (10-100 Days, Los Angeles, Vehicle Stop Black/white Rate Ratio Outcome) The x-axis is the number of days used in the data before and after the *BLM protest*. The y-axis is the unstandardized RDiT coefficient. Each facet characterizes a different kernel (Tri. = triangular, Uni. = uniform, Epa. = Epanechnikov) and polynomial (0, 1, 2, 3) specification. 95% CIs displayed from robust SEs.

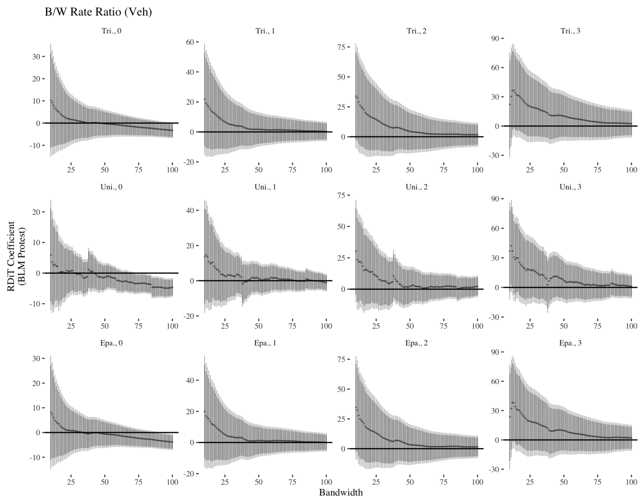

Figure F22: Alternative Bandwidths (10-100 Days, Philadelphia, Vehicle Stop Black/white Rate Ratio Outcome) The x-axis is the number of days used in the data before and after the *BLM protest*. The y-axis is the unstandardized RDiT coefficient. Each facet characterizes a different kernel (Tri. = triangular, Uni. = uniform, Epa. = Epanechnikov) and polynomial (0, 1, 2, 3) specification. 95% CIs displayed from robust SEs.

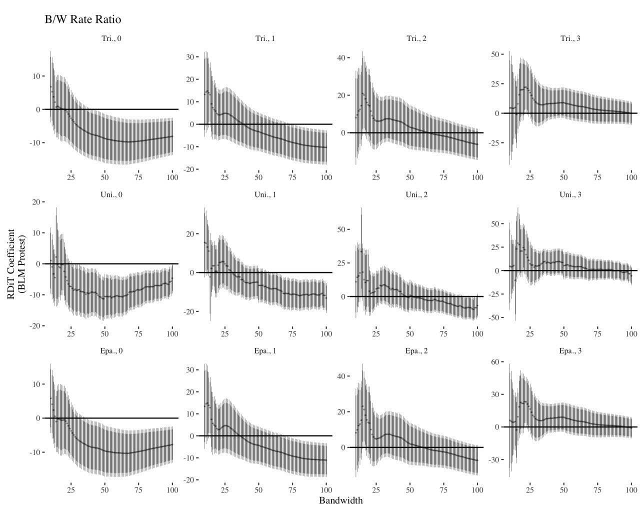

Figure F23: Alternative Bandwidths (10-100 Days, Seattle, Black/white Rate Ratio Outcome) The x-axis is the number of days used in the data before and after the *BLM protest*. The y-axis is the unstandardized RDiT coefficient. Each facet characterizes a different kernel (Tri. = triangular, Uni. = uniform, Epa. = Epanechnikov) and polynomial (0, 1, 2, 3) specification. 95% CIs displayed from robust SEs.

F.5 Crime

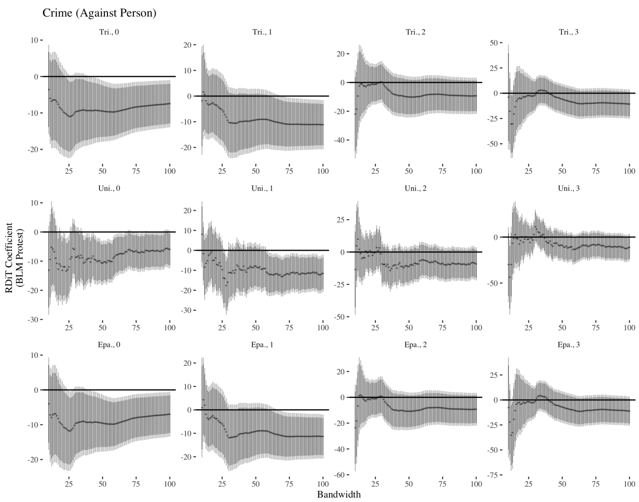

Figure F24: Alternative Bandwidths (10-100 Days, Austin, Crimes Against Person Outcome) The x-axis is the number of days used in the data before and after the *BLM protest*. The y-axis is the unstandardized RDiT coefficient. Each facet characterizes a different kernel (Tri. = triangular, Uni. = uniform, Epa. = Epanechnikov) and polynomial (0, 1, 2, 3) specification. 95% CIs displayed from robust SEs.

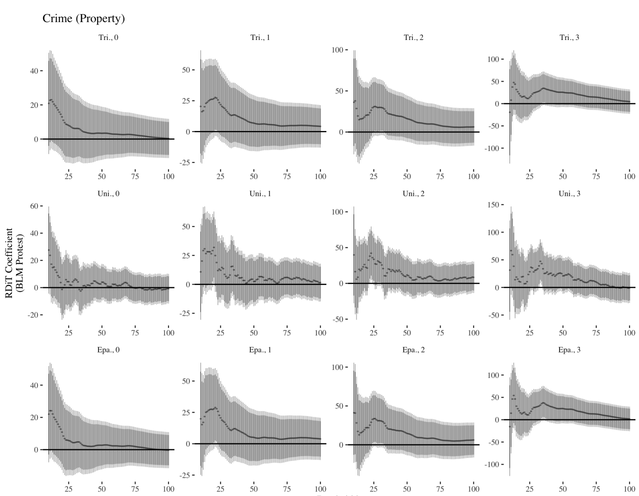

Figure F25: Alternative Bandwidths (10-100 Days, Austin, Crimes Against Property Outcome) The x-axis is the number of days used in the data before and after the *BLM protest*. The y-axis is the unstandardized RDiT coefficient. Each facet characterizes a different kernel (Tri. = triangular, Uni. = uniform, Epa. = Epanechnikov) and polynomial (0, 1, 2, 3) specification. 95% CIs displayed from robust SEs.

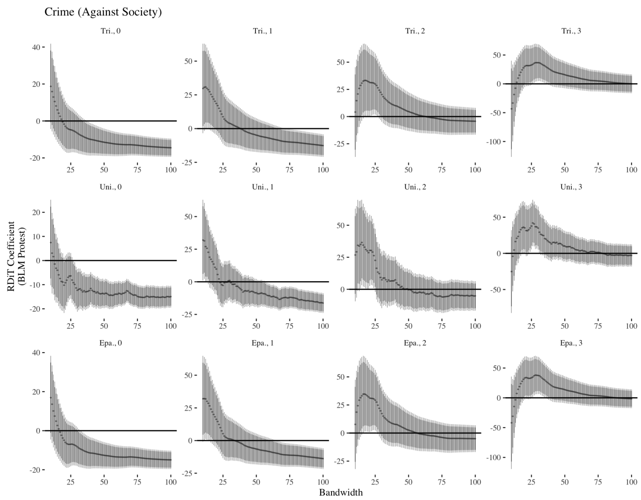

Figure F26: Alternative Bandwidths (10-100 Days, Austin, Crimes Against Society Outcome) The x-axis is the number of days used in the data before and after the *BLM protest*. The y-axis is the unstandardized RDiT coefficient. Each facet characterizes a different kernel (Tri. = triangular, Uni. = uniform, Epa. = Epanechnikov) and polynomial (0, 1, 2, 3) specification. 95% CIs displayed from robust SEs.

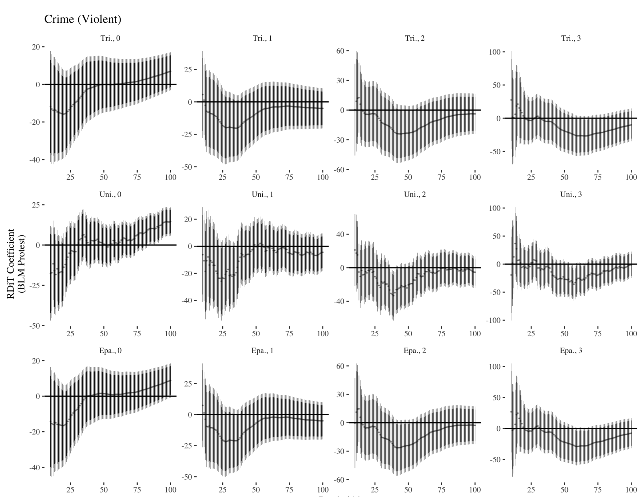

Figure F27: Alternative Bandwidths (10-100 Days, Los Angeles, Crimes Against Persons Outcome) The x-axis is the number of days used in the data before and after the *BLM protest*. The y-axis is the unstandardized RDiT coefficient. Each facet characterizes a different kernel (Tri. = triangular, Uni. = uniform, Epa. = Epanechnikov) and polynomial (0, 1, 2, 3) specification. 95% CIs displayed from robust SEs.

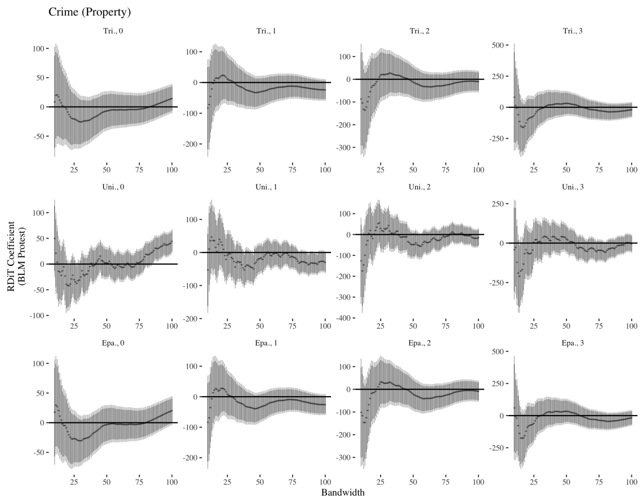

Figure F28: Alternative Bandwidths (10-100 Days, Los Angeles, Crimes Against Property Outcome) The x-axis is the number of days used in the data before and after the *BLM protest*. The y-axis is the unstandardized RDiT coefficient. Each facet characterizes a different kernel (Tri. = triangular, Uni. = uniform, Epa. = Epanechnikov) and polynomial (0, 1, 2, 3) specification. 95% CIs displayed from robust SEs.

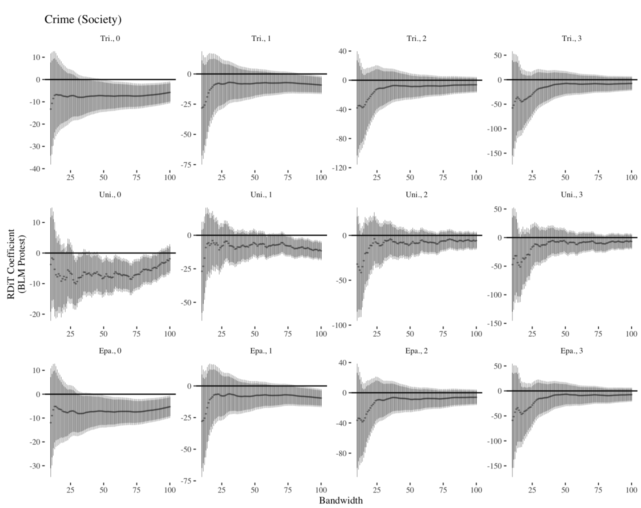

Figure F29: Alternative Bandwidths (10-100 Days, Los Angeles, Crimes Against Society Outcome) The x-axis is the number of days used in the data before and after the *BLM protest*. The y-axis is the unstandardized RDiT coefficient. Each facet characterizes a different kernel (Tri. = triangular, Uni. = uniform, Epa. = Epanechnikov) and polynomial (0, 1, 2, 3) specification. 95% CIs displayed from robust SEs.

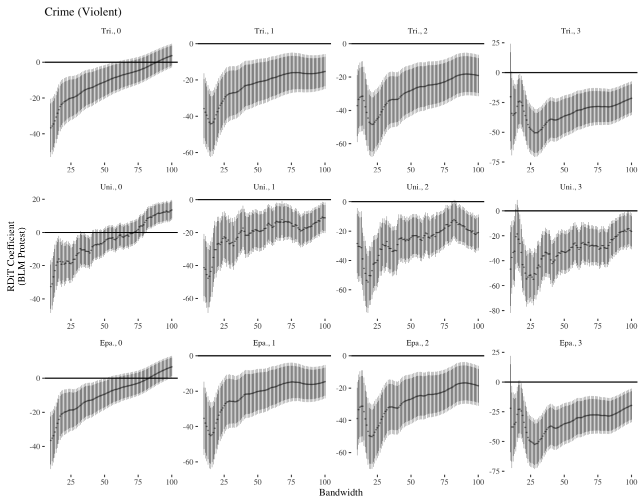

Figure F30: Alternative Bandwidths (10-100 Days, Philadelphia, Crimes Against Persons Outcome) The x-axis is the number of days used in the data before and after the *BLM protest*. The y-axis is the unstandardized RDiT coefficient. Each facet characterizes a different kernel (Tri. = triangular, Uni. = uniform, Epa. = Epanechnikov) and polynomial (0, 1, 2, 3) specification. 95% CIs displayed from robust SEs.

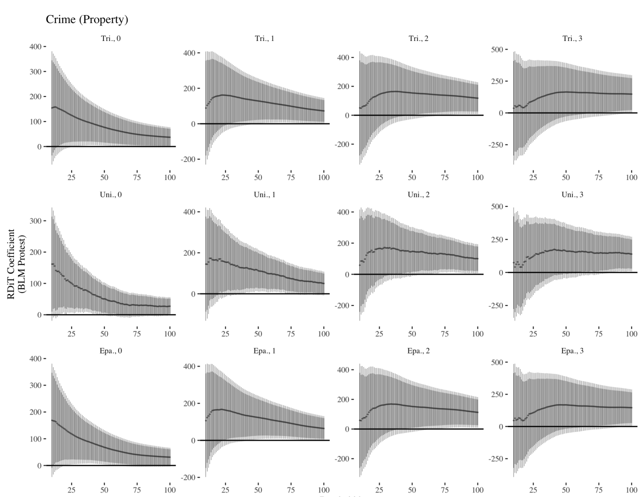

Figure F31: Alternative Bandwidths (10-100 Days, Philadelphia, Crimes Against Property Outcome) The x-axis is the number of days used in the data before and after the *BLM protest*. The y-axis is the unstandardized RDiT coefficient. Each facet characterizes a different kernel (Tri. = triangular, Uni. = uniform, Epa. = Epanechnikov) and polynomial (0, 1, 2, 3) specification. 95% CIs displayed from robust SEs.

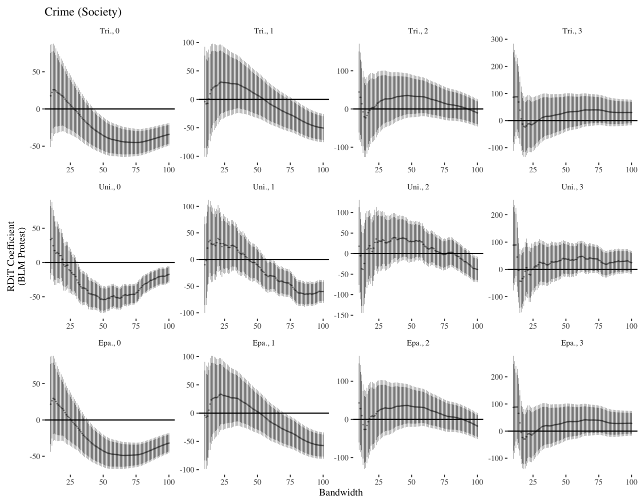

Figure F32: Alternative Bandwidths (10-100 Days, Philadelphia, Crimes Against Society Outcome) The x-axis is the number of days used in the data before and after the *BLM protest*. The y-axis is the unstandardized RDiT coefficient. Each facet characterizes a different kernel (Tri. = triangular, Uni. = uniform, Epa. = Epanechnikov) and polynomial (0, 1, 2, 3) specification. 95% CIs displayed from robust SEs.

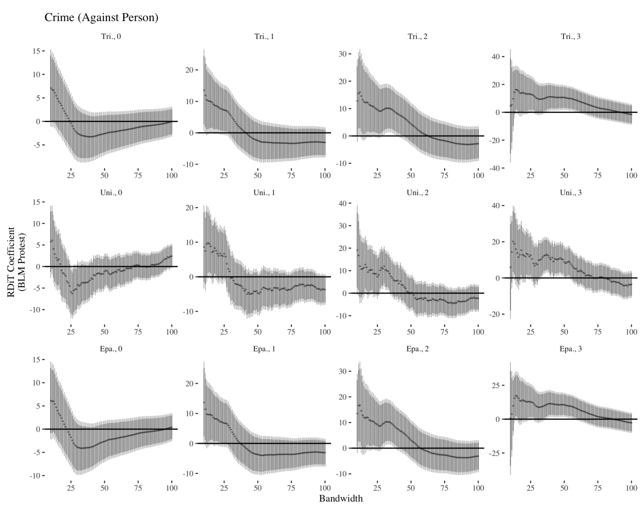

Figure F33: Alternative Bandwidths (10-100 Days, Seattle, Crimes Against Persons Outcome) The x-axis is the number of days used in the data before and after the *BLM protest*. The y-axis is the unstandardized RDiT coefficient. Each facet characterizes a different kernel (Tri. = triangular, Uni. = uniform, Epa. = Epanechnikov) and polynomial (0, 1, 2, 3) specification. 95% CIs displayed from robust SEs.

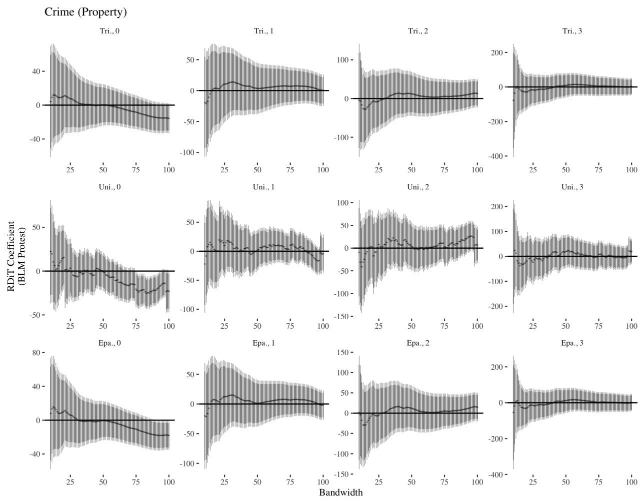

Figure F34: Alternative Bandwidths (10-100 Days, Seattle, Crimes Against Property Outcome) The x-axis is the number of days used in the data before and after the *BLM protest*. The y-axis is the unstandardized RDiT coefficient. Each facet characterizes a different kernel (Tri. = triangular, Uni. = uniform, Epa. = Epanechnikov) and polynomial (0, 1, 2, 3) specification. 95% CIs displayed from robust SEs.

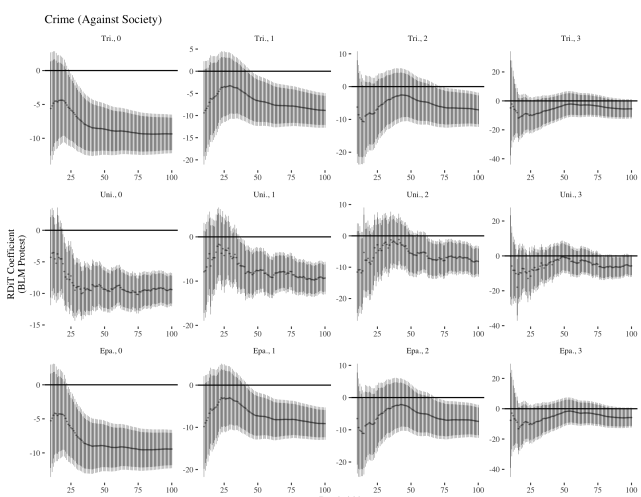

Figure F35: Alternative Bandwidths (10-100 Days, Seattle, Crimes Against Society Outcome) The x-axis is the number of days used in the data before and after the *BLM protest*. The y-axis is the unstandardized RDiT coefficient. Each facet characterizes a different kernel (Tri. = triangular, Uni. = uniform, Epa. = Epanechnikov) and polynomial (0, 1, 2, 3) specification. 95% CIs displayed from robust SEs.

G Long-Term Effects

G.1 Policing Activity

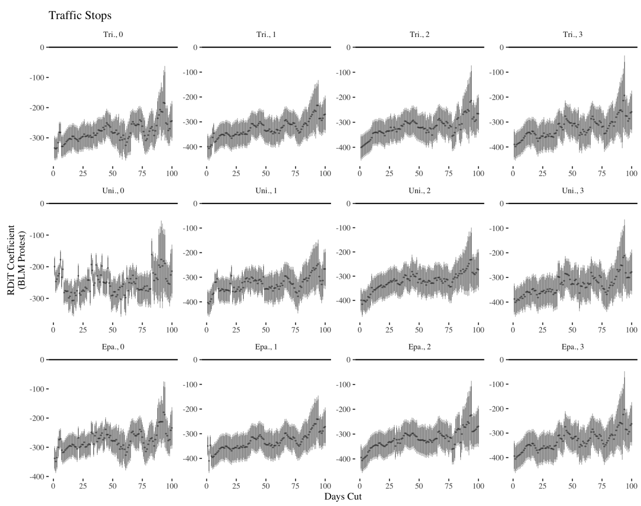

Figure G36: Assessing Persistence of Decreases in Policing Activity post-BLM Protest (Austin, Traffic Stop Outcome) X-axis is the number of days cut from the time series post-*BLM protest*. Y-axis is the unstandardized RDiT coefficient characterizing the discontinuous effect in the outcome between the time period immediately before the BLM protests and however many days after the onset of the BLM protests.

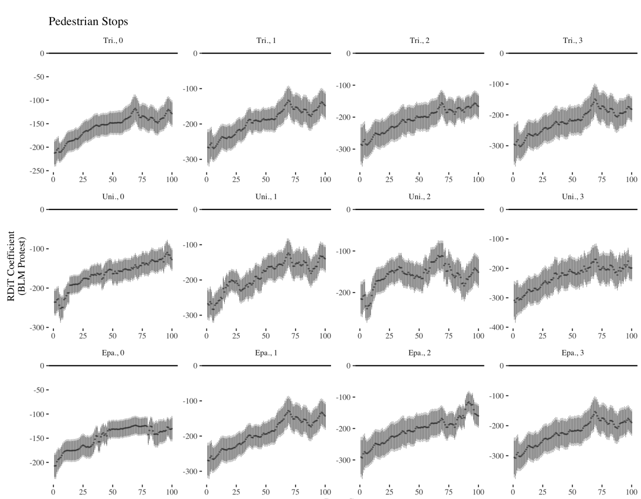

Figure G37: Assessing Persistence of Decreases in Policing Activity post-BLM Protest (Los Angeles, Pedestrian Stop Outcome) X-axis is the number of days cut from the time series post-*BLM protest*. Y-axis is the unstandardized RDiT coefficient characterizing the discontinuous effect in the outcome between the time period immediately before the BLM protests and however many days after the onset of the BLM protests.

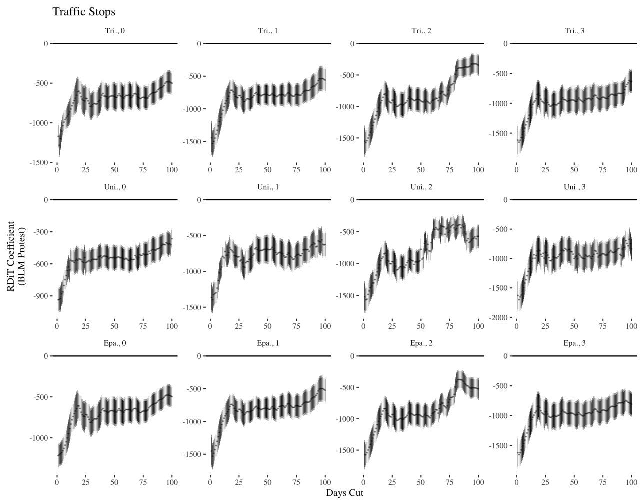

Figure G38: Assessing Persistence of Decreases in Policing Activity post-BLM Protest (Los Angeles, Traffic Stop Outcome) X-axis is the number of days cut from the time series post-*BLM protest*. Y-axis is the unstandardized RDiT coefficient characterizing the discontinuous effect in the outcome between the time period immediately before the BLM protests and however many days after the onset of the BLM protests.

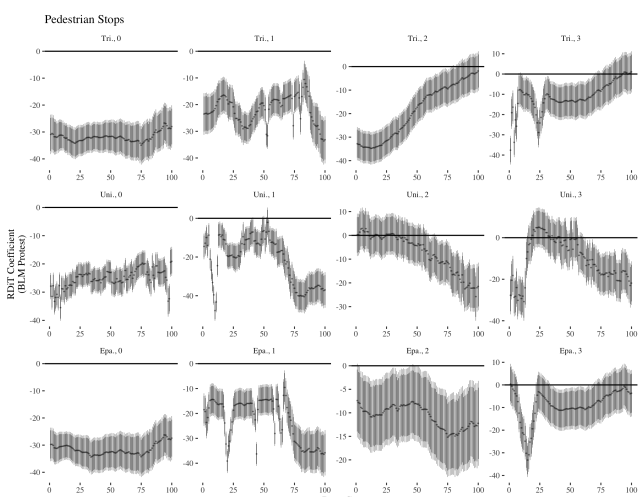

Figure G39: Assessing Persistence of Decreases in Policing Activity post-BLM Protest (Philadelphia, Pedestrian Stop Outcome) X-axis is the number of days cut from the time series post-*BLM protest*. Y-axis is the unstandardized RDiT coefficient characterizing the discontinuous effect in the outcome between the time period immediately before the BLM protests and however many days after the onset of the BLM protests.

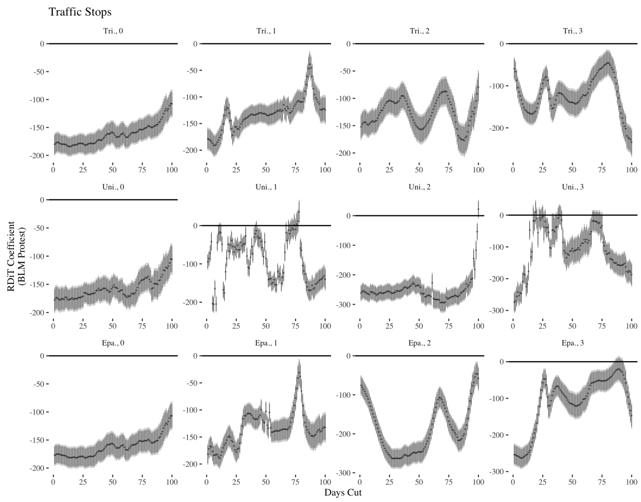

Figure G40: Assessing Persistence of Decreases in Policing Activity post-BLM Protest (Philadelphia, Traffic Stop Outcome) X-axis is the number of days cut from the time series post-*BLM protest*. Y-axis is the unstandardized RDiT coefficient characterizing the discontinuous effect in the outcome between the time period immediately before the BLM protests and however many days after the onset of the BLM protests.

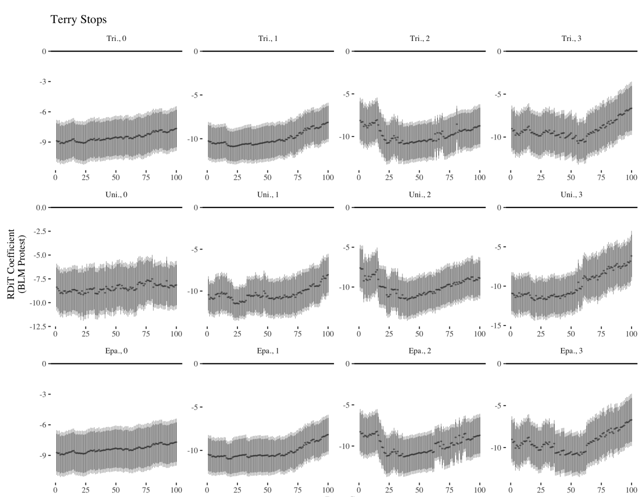

Figure G41: Assessing Persistence of Decreases in Policing Activity post-BLM Protest (Seattle, Terry Stop Outcome) X-axis is the number of days cut from the time series post-*BLM protest*. Y-axis is the unstandardized RDiT coefficient characterizing the discontinuous effect in the outcome between the time period immediately before the BLM protests and however many days after the onset of the BLM protests.

G.2 Hit Rates

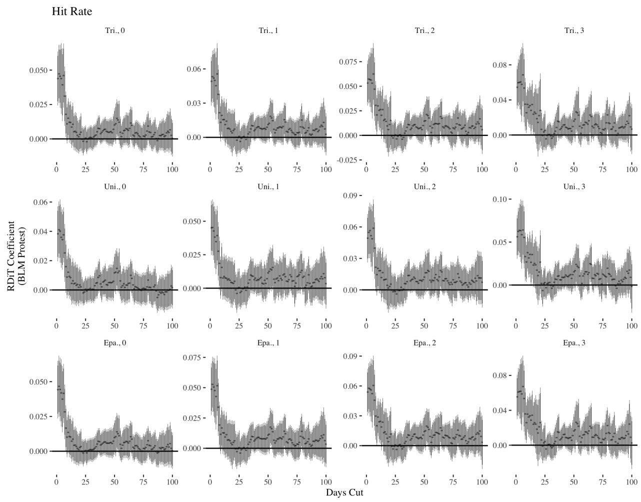

Figure G42: Assessing Persistence of RDiT Effects post-BLM Protest (Austin, Hit Rate Outcome) X-axis is the number of days cut from the time series post-*BLM protest*. Y-axis is the unstandardized RDiT coefficient characterizing the discontinuous effect in the outcome between the time period immediately before the BLM protests and however many days after the onset of the BLM protests.

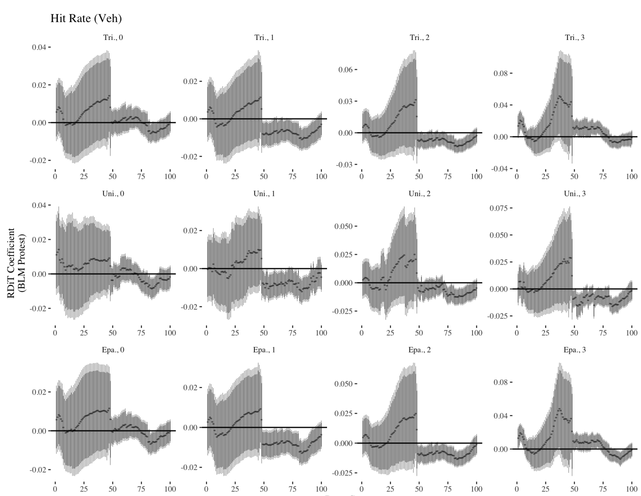

Figure G43: Assessing Persistence of RDiT Effects post-BLM Protest (Los Angeles, Traffic Hit Rate Outcome) X-axis is the number of days cut from the time series post-*BLM protest*. Y-axis is the unstandardized RDiT coefficient characterizing the discontinuous effect in the outcome between the time period immediately before the BLM protests and however many days after the onset of the BLM protests.

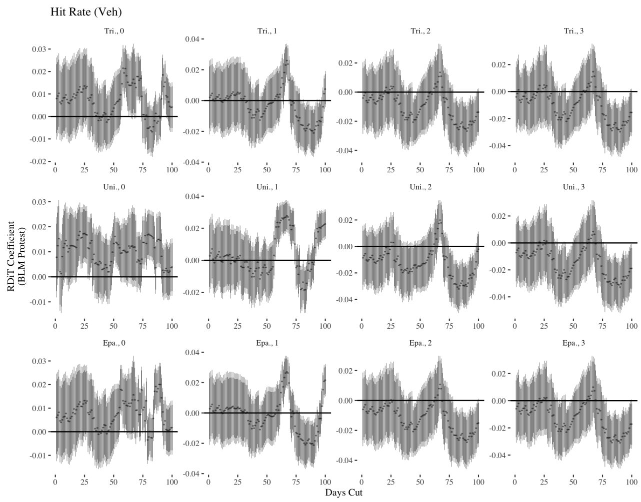

Figure G44: Assessing Persistence of RDiT Effects post-BLM Protest (Philadelphia, Traffic Hit Rate Outcome) X-axis is the number of days cut from the time series post-*BLM protest*. Y-axis is the unstandardized RDiT coefficient characterizing the discontinuous effect in the outcome between the time period immediately before the BLM protests and however many days after the onset of the BLM protests.

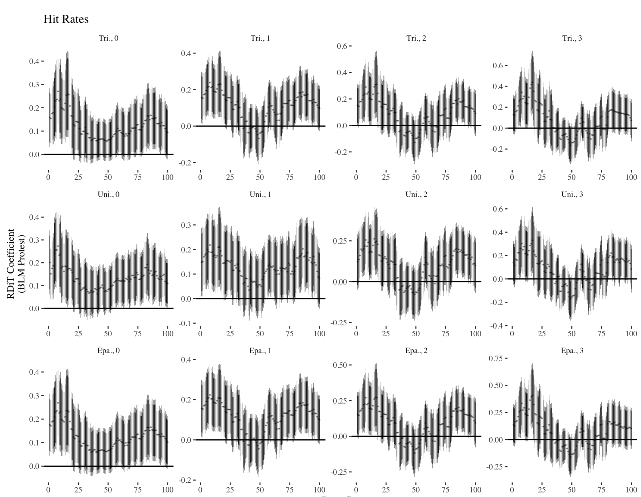

Figure G45: Assessing Persistence of RDiT Effects post-BLM Protest (Seattle, Terry Hit Rate Outcome) X-axis is the number of days cut from the time series post-*BLM protest*. Y-axis is the unstandardized RDiT coefficient characterizing the discontinuous effect in the outcome between the time period immediately before the BLM protests and however many days after the onset of the BLM protests.

G.3 Arrest Rates

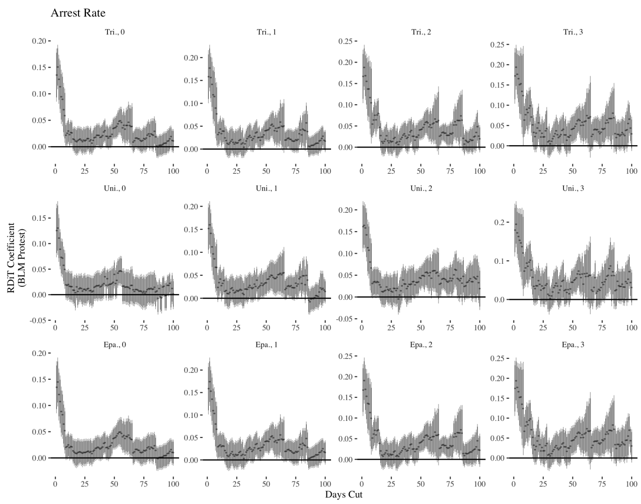

Figure G46: Assessing Persistence of RDiT Effects post-BLM Protest (Austin, Arrest Rate Outcome) X-axis is the number of days cut from the time series post-*BLM protest*. Y-axis is the unstandardized RDiT coefficient characterizing the discontinuous effect in the outcome between the time period immediately before the BLM protests and however many days after the onset of the BLM protests.

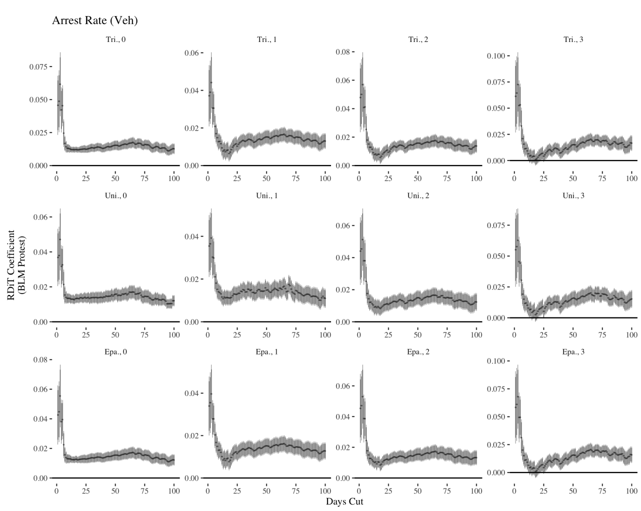

Figure G47: Assessing Persistence of RDiT Effects post-BLM Protest (Los Angeles, Traffic Arrest Rate Outcome) X-axis is the number of days cut from the time series post-*BLM protest*. Y-axis is the unstandardized RDiT coefficient characterizing the discontinuous effect in the outcome between the time period immediately before the BLM protests and however many days after the onset of the BLM protests.

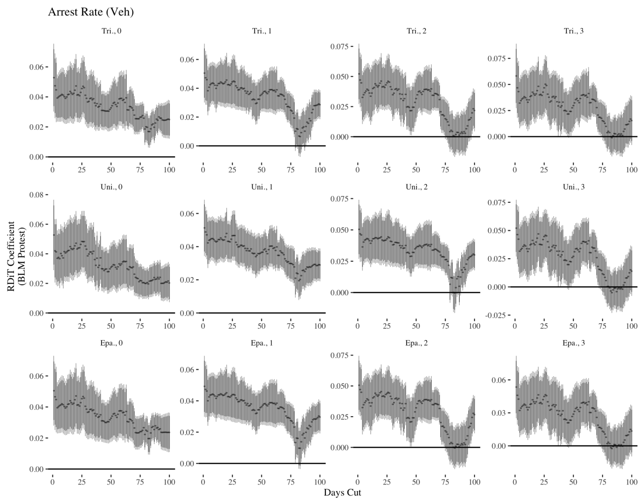

Figure G48: Assessing Persistence of RDiT Effects post-BLM Protest (Philadelphia, Traffic Arrest Rate Outcome) X-axis is the number of days cut from the time series post-*BLM protest*. Y-axis is the unstandardized RDiT coefficient characterizing the discontinuous effect in the outcome between the time period immediately before the BLM protests and however many days after the onset of the BLM protests.

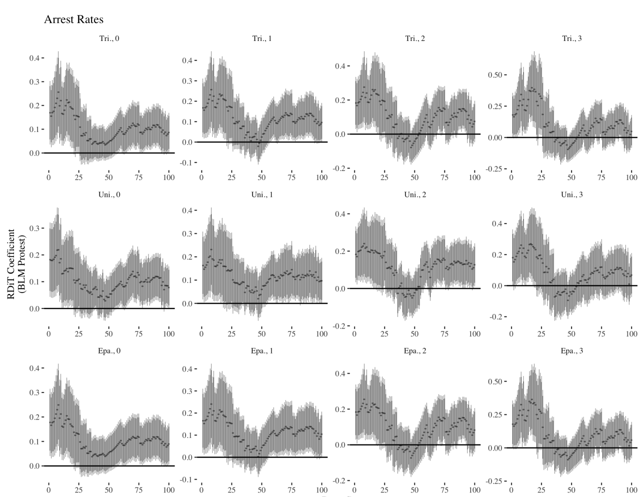

Figure G49: Assessing Persistence of RDiT Effects post-BLM Protest (Seattle, Terry Arrest Rate Outcome) X-axis is the number of days cut from the time series post-*BLM protest*. Y-axis is the unstandardized RDiT coefficient characterizing the discontinuous effect in the outcome between the time period immediately before the BLM protests and however many days after the onset of the BLM protests.

G.4 Rate Ratios

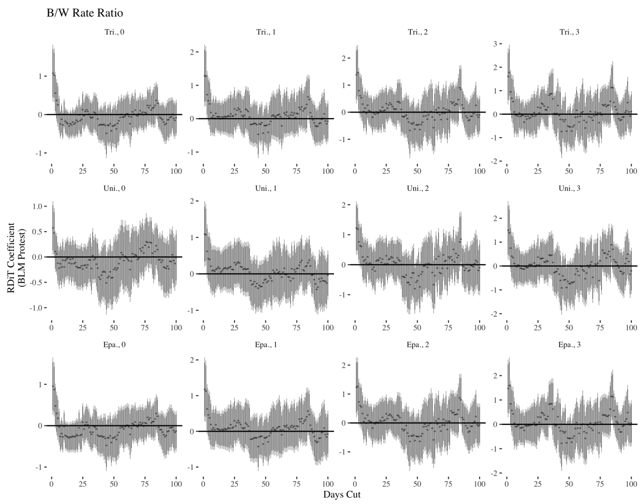

Figure G50: Assessing Persistence of RDiT Effects post-BLM Protest (Austin, Black/white Traffic Stop Rate Ratio Outcome) X-axis is the number of days cut from the time series post-*BLM protest*. Y-axis is the unstandardized RDiT coefficient characterizing the discontinuous effect in the outcome between the time period immediately before the BLM protests and however many days after the onset of the BLM protests.

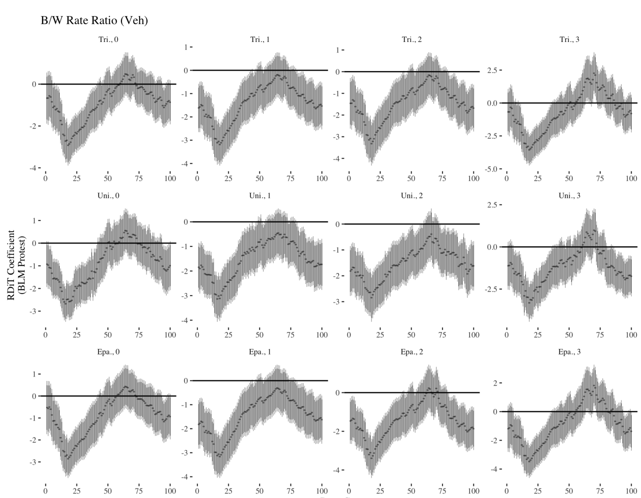

Figure G51: Assessing Persistence of RDiT Effects post-BLM Protest (Los Angeles, Black/white Vehicle Stop Rate Ratio Outcome) X-axis is the number of days cut from the time series post-*BLM protest*. Y-axis is the unstandardized RDiT coefficient characterizing the discontinuous effect in the outcome between the time period immediately before the BLM protests and however many days after the onset of the BLM protests.

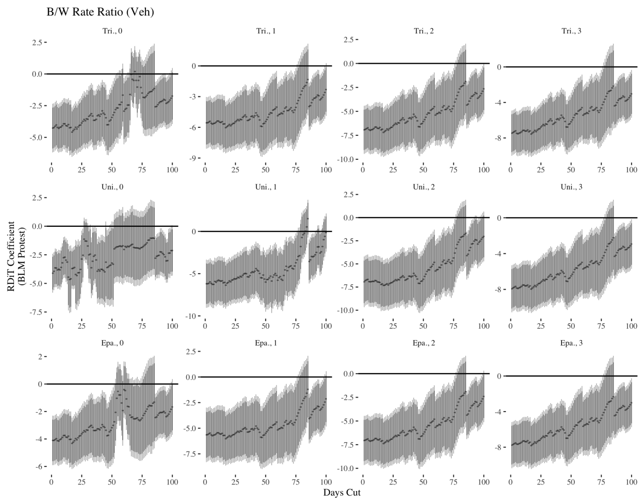

Figure G52: Assessing Persistence of RDiT Effects post-BLM Protest (Philadelphia, Black/white Traffic Stop Rate Ratio Outcome) X-axis is the number of days cut from the time series post-*BLM protest*. Y-axis is the unstandardized RDiT coefficient characterizing the discontinuous effect in the outcome between the time period immediately before the BLM protests and however many days after the onset of the BLM protests.

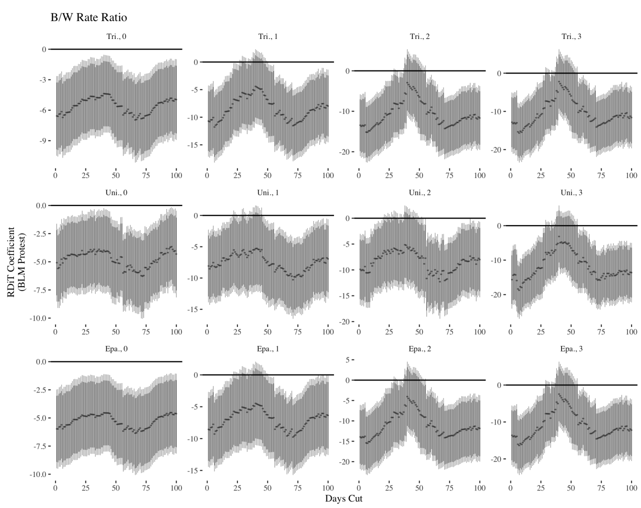

Figure G53: Assessing Persistence of RDiT Effects post-BLM Protest (Seattle, Black/white Terry Stop Rate Ratio Outcome) X-axis is the number of days cut from the time series post-*BLM protest*. Y-axis is the unstandardized RDiT coefficient characterizing the discontinuous effect in the outcome between the time period immediately before the BLM protests and however many days after the onset of the BLM protests.

G.5 Crime

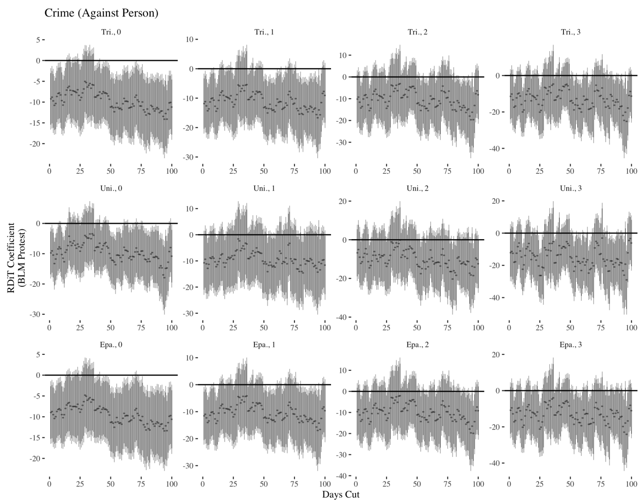

Figure G54: Assessing Persistence of RDiT Effects post-BLM Protest (Austin, Crimes Against Persons Outcome) X-axis is the number of days cut from the time series post-*BLM protest*. Y-axis is the unstandardized RDiT coefficient characterizing the discontinuous effect in the outcome between the time period immediately before the BLM protests and however many days after the onset of the BLM protests.

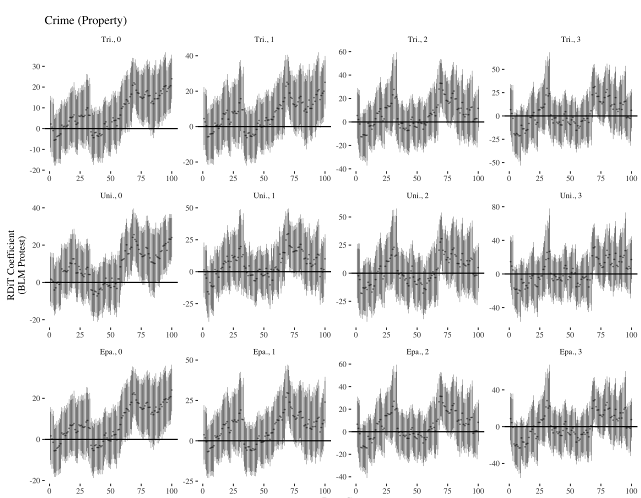

Figure G55: Assessing Persistence of RDiT Effects post-BLM Protest (Austin, Crimes Against Property Outcome) X-axis is the number of days cut from the time series post-*BLM protest*. Y-axis is the unstandardized RDiT coefficient characterizing the discontinuous effect in the outcome between the time period immediately before the BLM protests and however many days after the onset of the BLM protests.

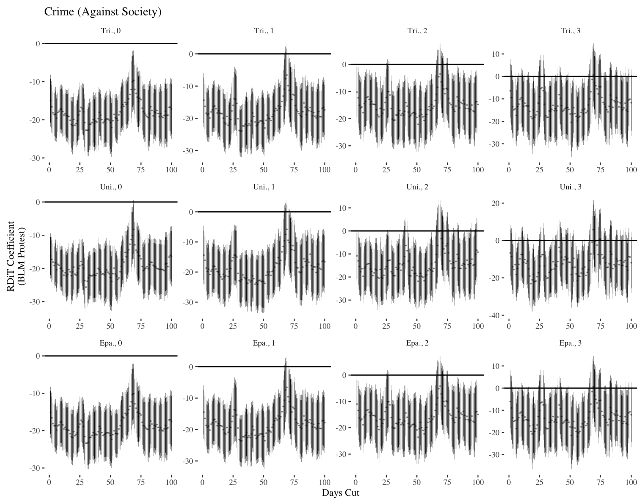

Figure G56: Assessing Persistence of RDiT Effects post-BLM Protest (Austin, Crimes Against Society Outcome) X-axis is the number of days cut from the time series post-*BLM protest*. Y-axis is the unstandardized RDiT coefficient characterizing the discontinuous effect in the outcome between the time period immediately before the BLM protests and however many days after the onset of the BLM protests.

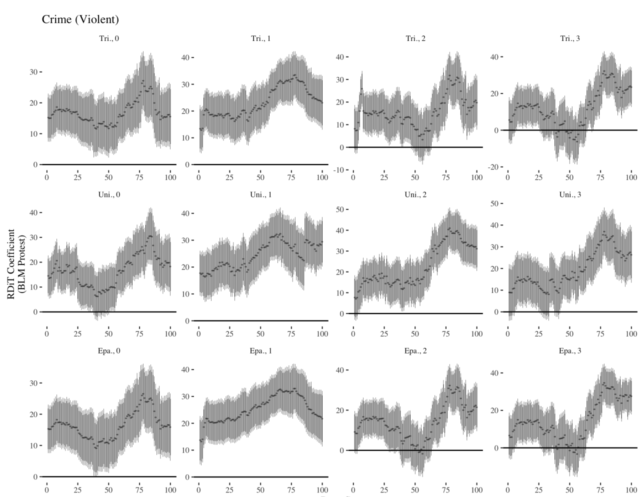

Figure G57: Assessing Persistence of RDiT Effects post-BLM Protest (Los Angeles, Crimes Against Person Outcome) X-axis is the number of days cut from the time series post-*BLM protest*. Y-axis is the unstandardized RDiT coefficient characterizing the discontinuous effect in the outcome between the time period immediately before the BLM protests and however many days after the onset of the BLM protests.

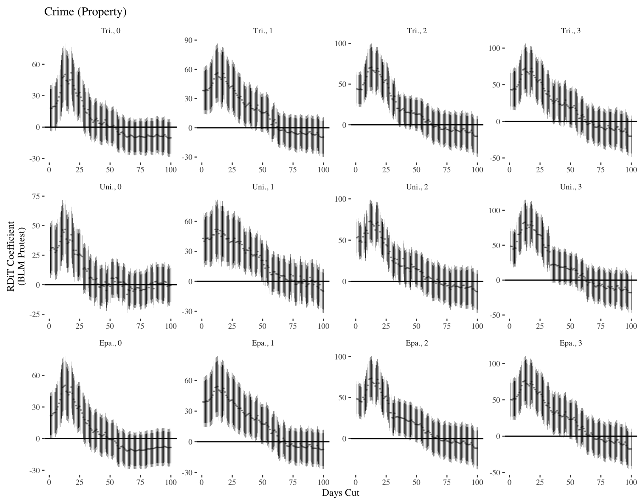

Figure G58: Assessing Persistence of RDiT Effects post-BLM Protest (Los Angeles, Crimes Against Property Outcome) X-axis is the number of days cut from the time series post-*BLM protest*. Y-axis is the unstandardized RDiT coefficient characterizing the discontinuous effect in the outcome between the time period immediately before the BLM protests and however many days after the onset of the BLM protests.

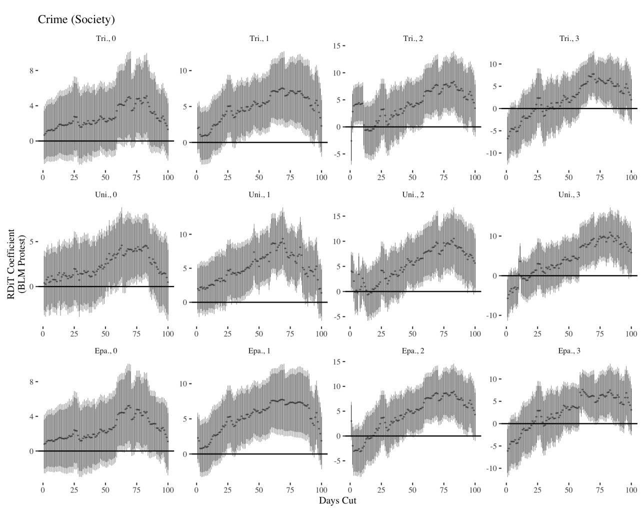

Figure G59: Assessing Persistence of RDiT Effects post-BLM Protest (Los Angeles, Crimes Against Society Outcome) X-axis is the number of days cut from the time series post-*BLM protest*. Y-axis is the unstandardized RDiT coefficient characterizing the discontinuous effect in the outcome between the time period immediately before the BLM protests and however many days after the onset of the BLM protests.

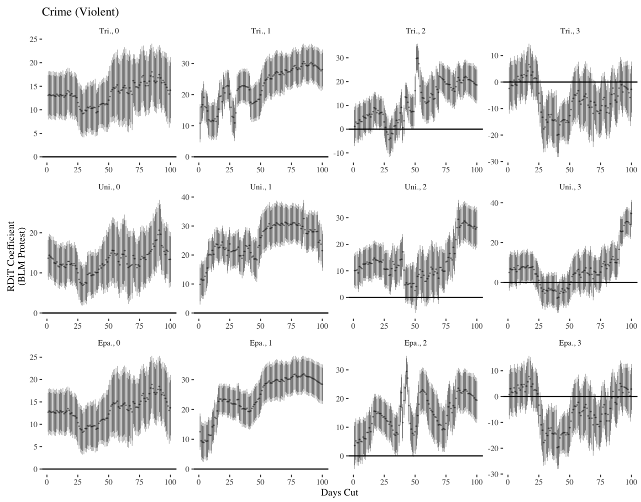

Figure G60: Assessing Persistence of RDiT Effects post-BLM Protest (Philadelphia, Crimes Against Persons Outcome) X-axis is the number of days cut from the time series post-*BLM protest*. Y-axis is the unstandardized RDiT coefficient characterizing the discontinuous effect in the outcome between the time period immediately before the BLM protests and however many days after the onset of the BLM protests.

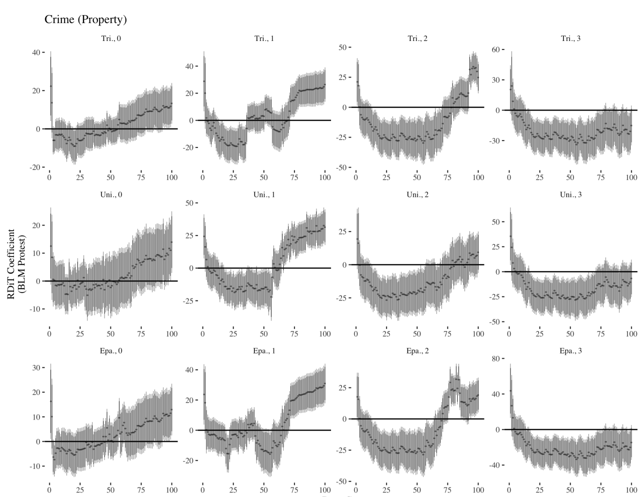

Figure G61: Assessing Persistence of RDiT Effects post-BLM Protest (Philadelphia, Crimes Against Property Outcome) X-axis is the number of days cut from the time series post-*BLM protest*. Y-axis is the unstandardized RDiT coefficient characterizing the discontinuous effect in the outcome between the time period immediately before the BLM protests and however many days after the onset of the BLM protests.

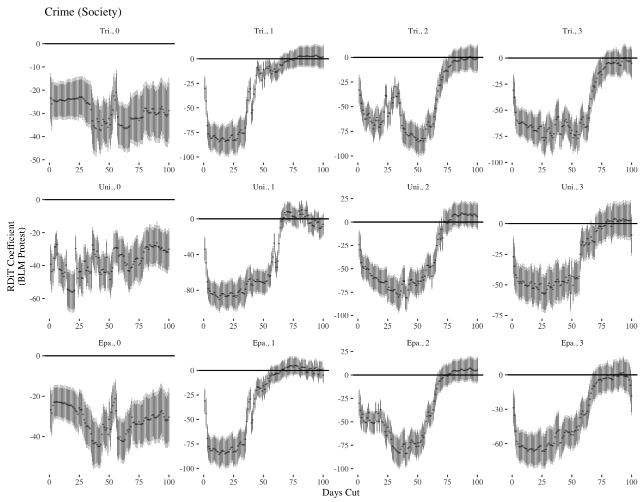

Figure G62: Assessing Persistence of RDiT Effects post-BLM Protest (Philadelphia, Crimes Against Society Outcome) X-axis is the number of days cut from the time series post-*BLM protest*. Y-axis is the unstandardized RDiT coefficient characterizing the discontinuous effect in the outcome between the time period immediately before the BLM protests and however many days after the onset of the BLM protests.

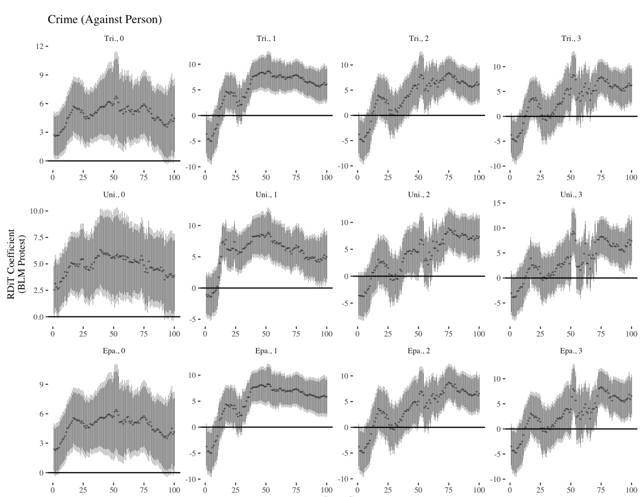

Figure G63: Assessing Persistence of RDiT Effects post-BLM Protest (Seattle, Crimes Against Persons Outcome) X-axis is the number of days cut from the time series post-*BLM protest*. Y-axis is the unstandardized RDiT coefficient characterizing the discontinuous effect in the outcome between the time period immediately before the BLM protests and however many days after the onset of the BLM protests.

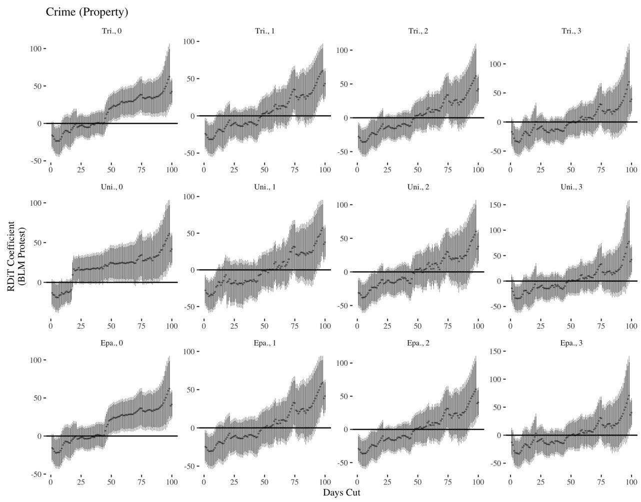

Figure G64: Assessing Persistence of RDiT Effects post-BLM Protest (Seattle, Crimes Against Property Outcome) X-axis is the number of days cut from the time series post-*BLM protest*. Y-axis is the unstandardized RDiT coefficient characterizing the discontinuous effect in the outcome between the time period immediately before the BLM protests and however many days after the onset of the BLM protests.

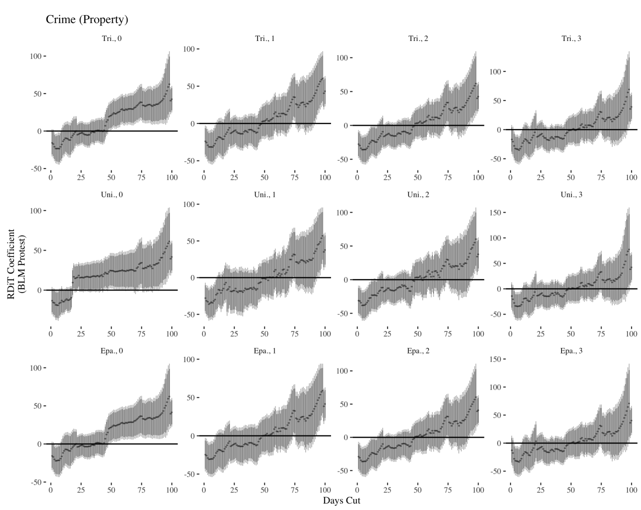

Figure G65: Assessing Persistence of RDiT Effects post-BLM Protest (Seattle, Crimes Against Society Outcome) X-axis is the number of days cut from the time series post-*BLM protest*. Y-axis is the unstandardized RDiT coefficient characterizing the discontinuous effect in the outcome between the time period immediately before the BLM protests and however many days after the onset of the BLM protests.

H Seasonal Placebo Test

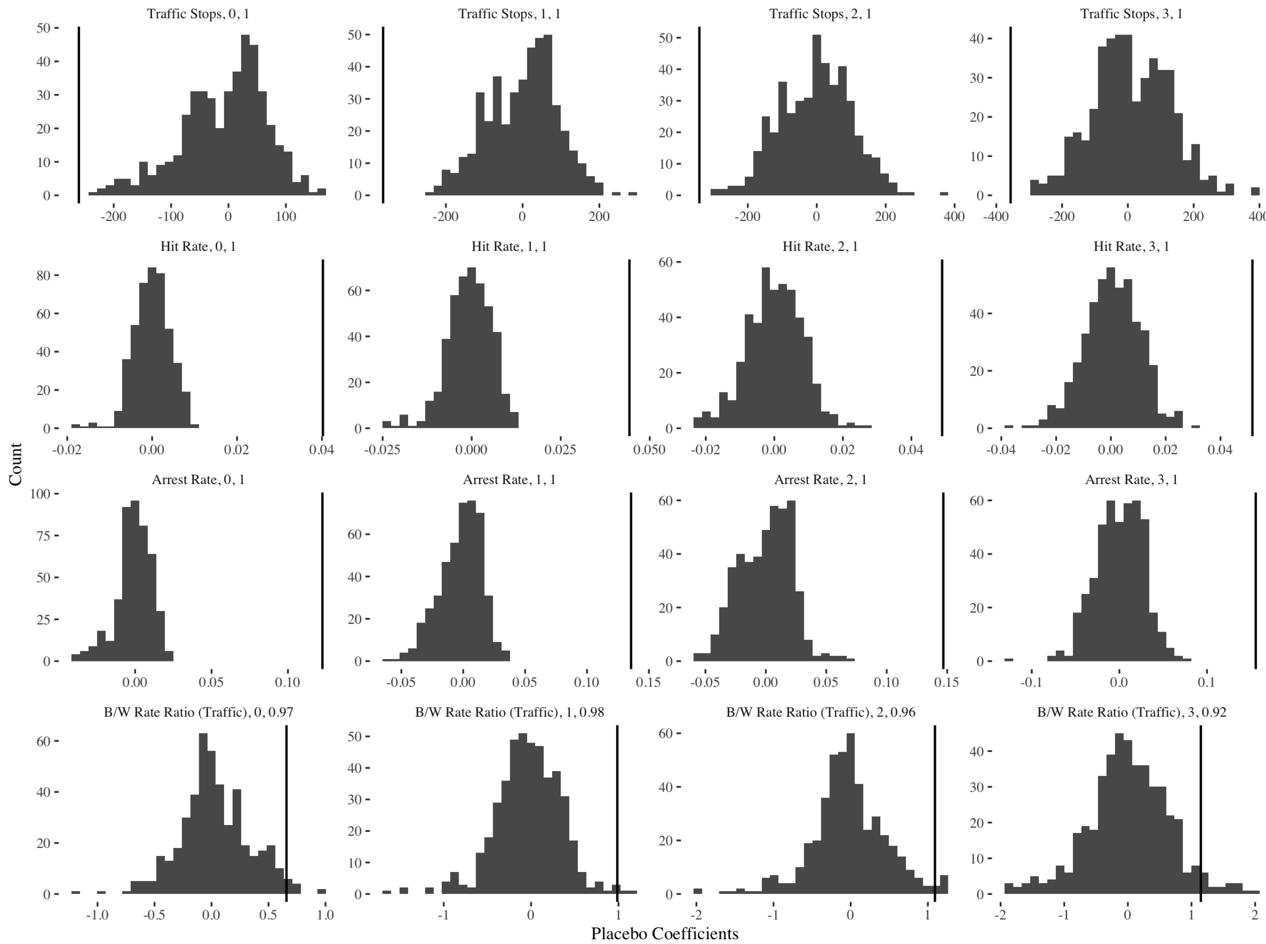

Figure H66: Temporal Placebo Tests Comparing True Post-BLM Discontinuous Effect to Pre-BLM Discontinuous Effects (Austin). The x-axis characterizes pre-BLM placebo coefficients. Solid vertical line denotes true post-BLM protest coefficient. Each facet denotes an outcome, polynomial degree, and the proportion of placebo coefficients (converted to absolute value) that the true coefficient (converted to absolute value) is larger than.

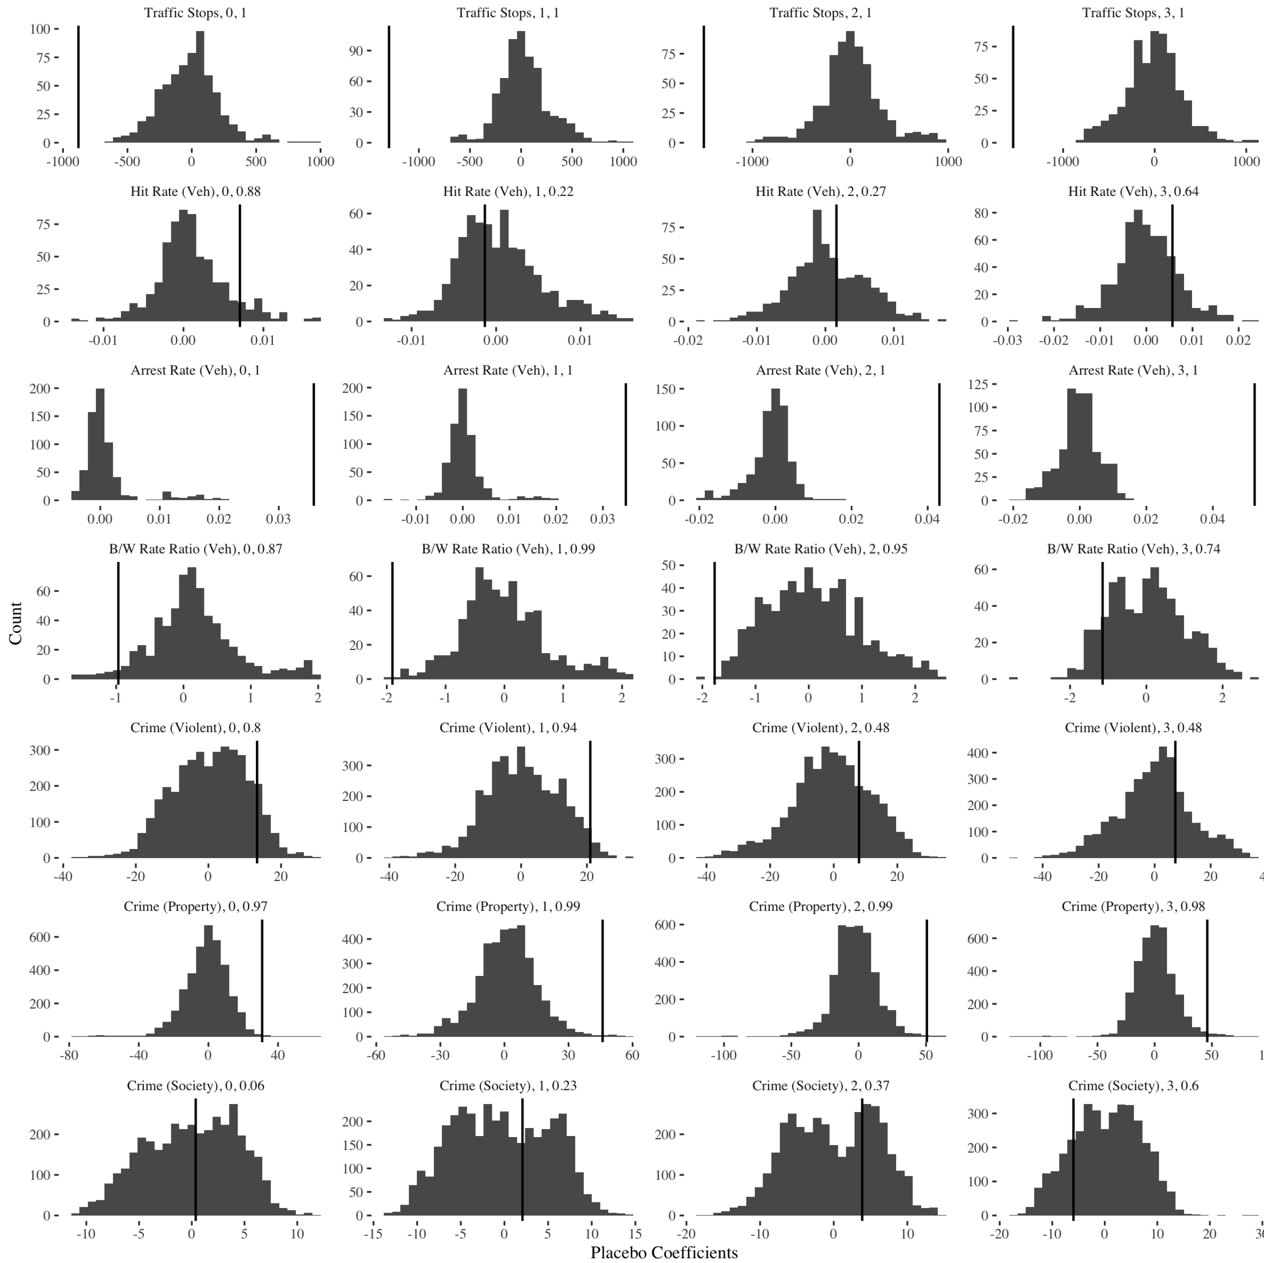

Figure H67: Temporal Placebo Tests Comparing True Post-BLM Discontinuous Effect to Pre-BLM Discontinuous Effects (Los Angeles). The x-axis characterizes pre-BLM placebo coefficients. Solid vertical line denotes true post-BLM protest coefficient. Each facet denotes an outcome, polynomial degree, and the proportion of placebo coefficients (converted to absolute value) that the true coefficient (converted to absolute value) is larger than.

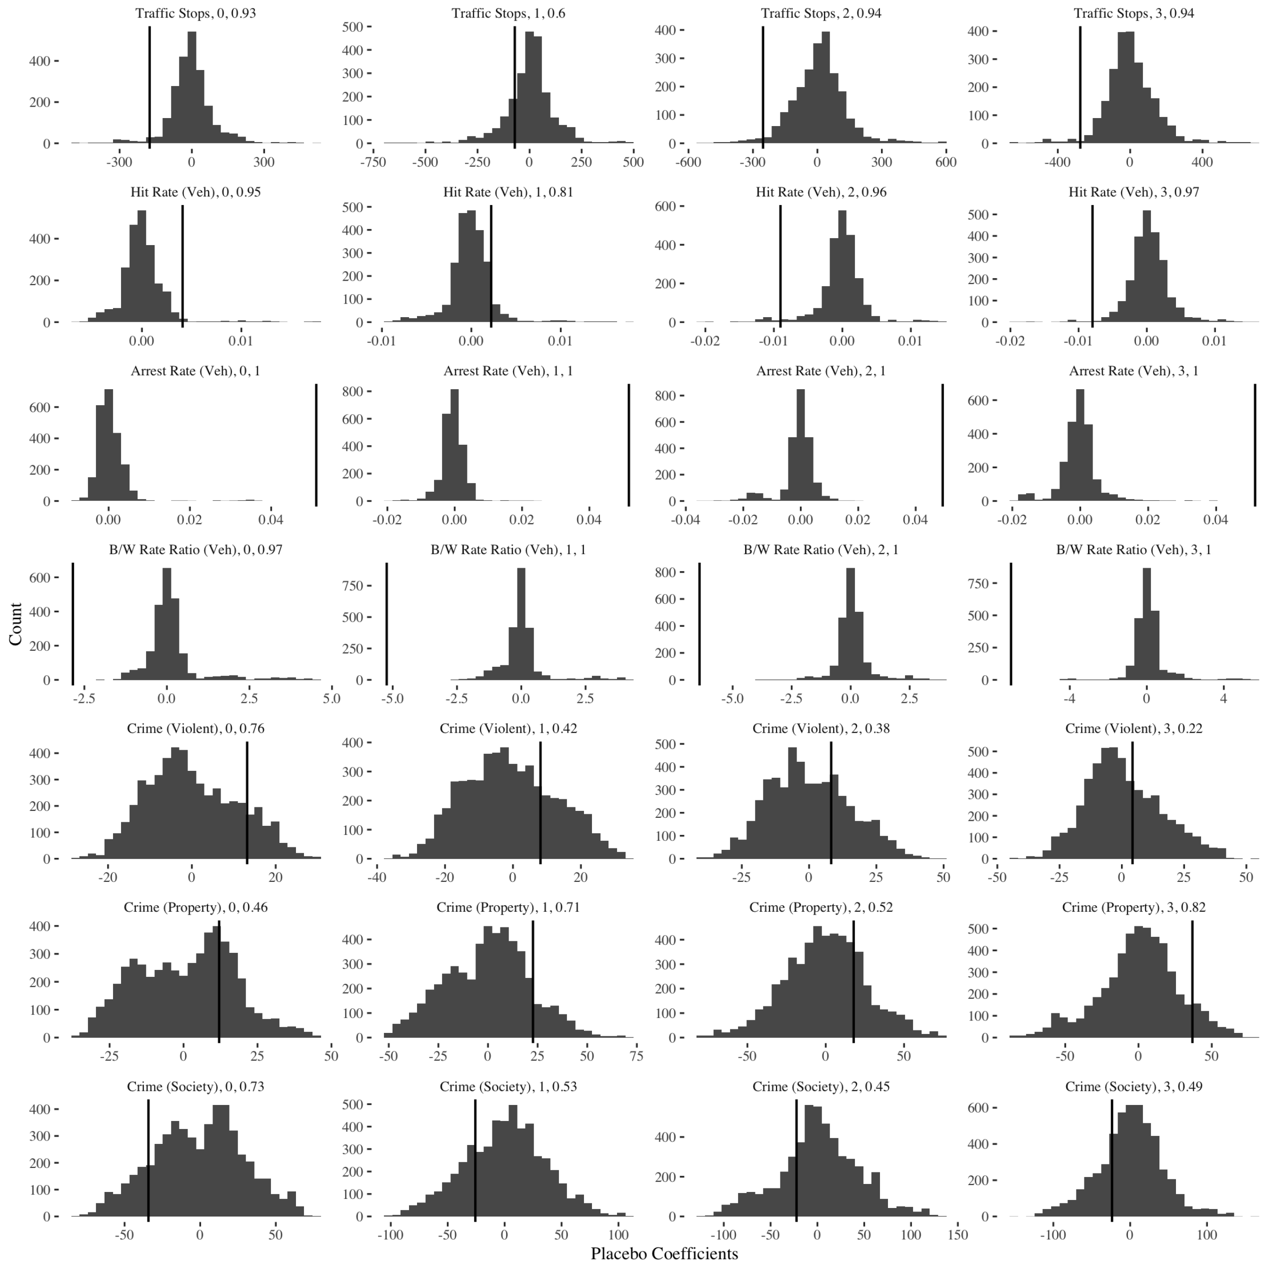

Figure H68: Temporal Placebo Tests Comparing True Post-BLM Discontinuous Effect to Pre-BLM Discontinuous Effects (Philadelphia). The x-axis characterizes pre-BLM placebo coefficients. Solid vertical line denotes true post-BLM protest coefficient. Each facet denotes an outcome, polynomial degree, and the proportion of placebo coefficients (converted to absolute value) that the true coefficient (converted to absolute value) is larger than.

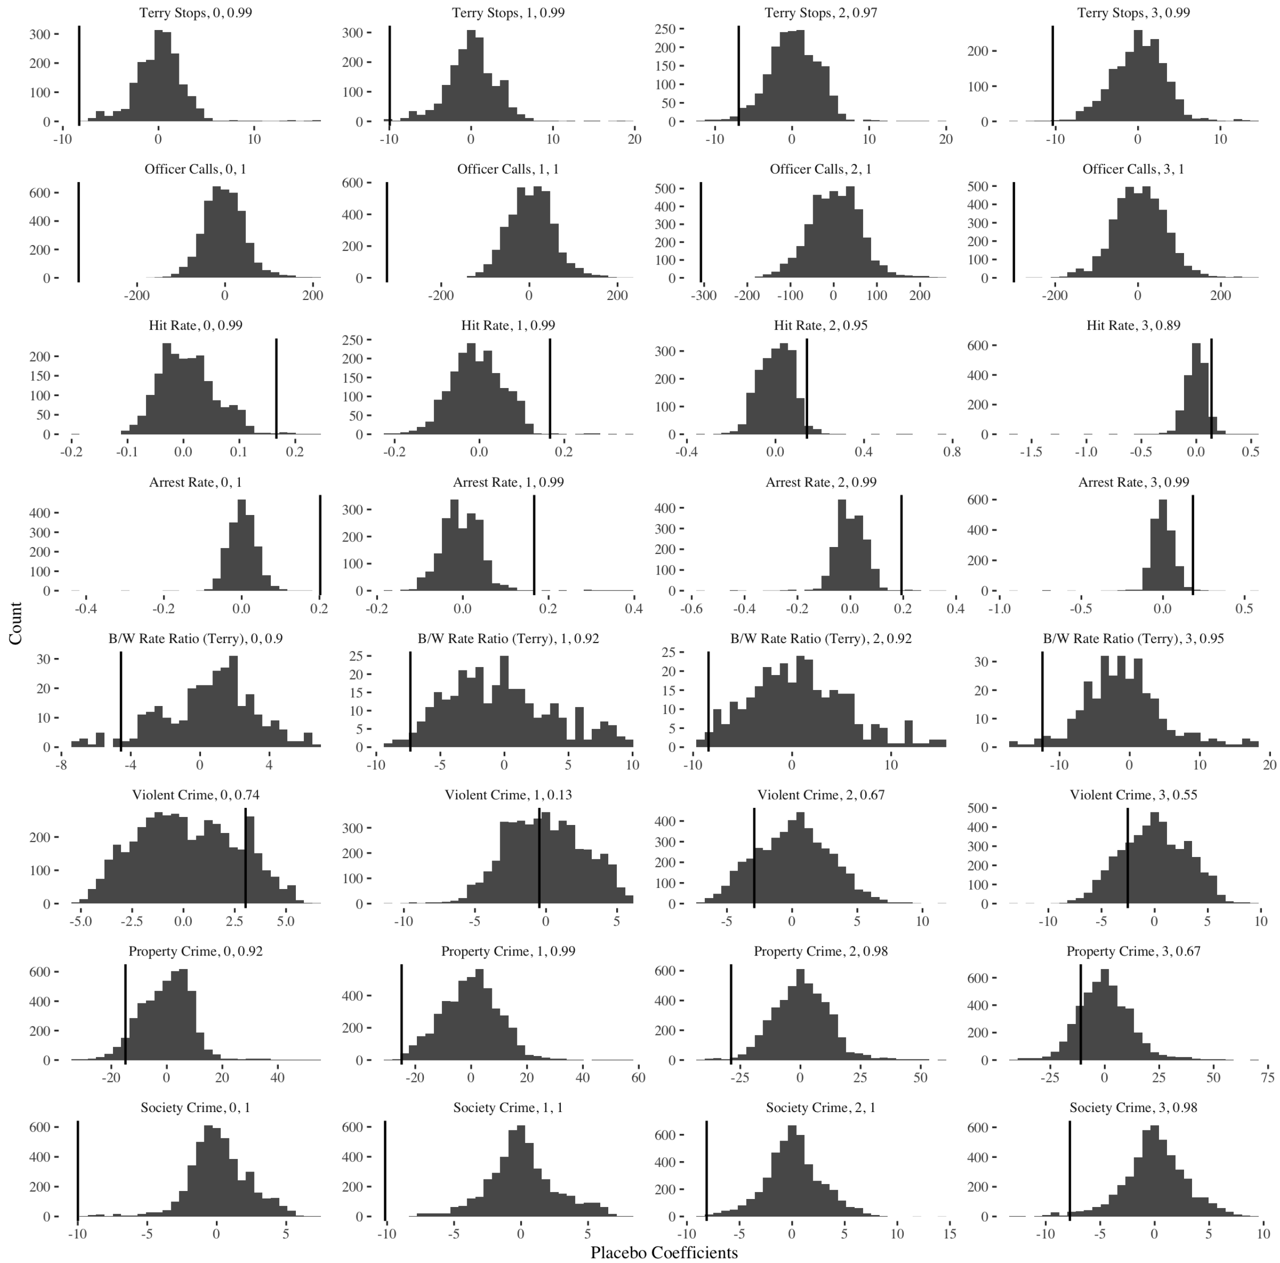

Figure H69: Temporal Placebo Tests Comparing True Post-BLM Discontinuous Effect to Pre-BLM Discontinuous Effects (Seattle). The x-axis characterizes pre-BLM placebo coefficients. Solid vertical line denotes true post-BLM protest coefficient. Each facet denotes an outcome, polynomial degree, and the proportion of placebo coefficients (converted to absolute value) that the true coefficient (converted to absolute value) is larger than.

I San Diego Replication

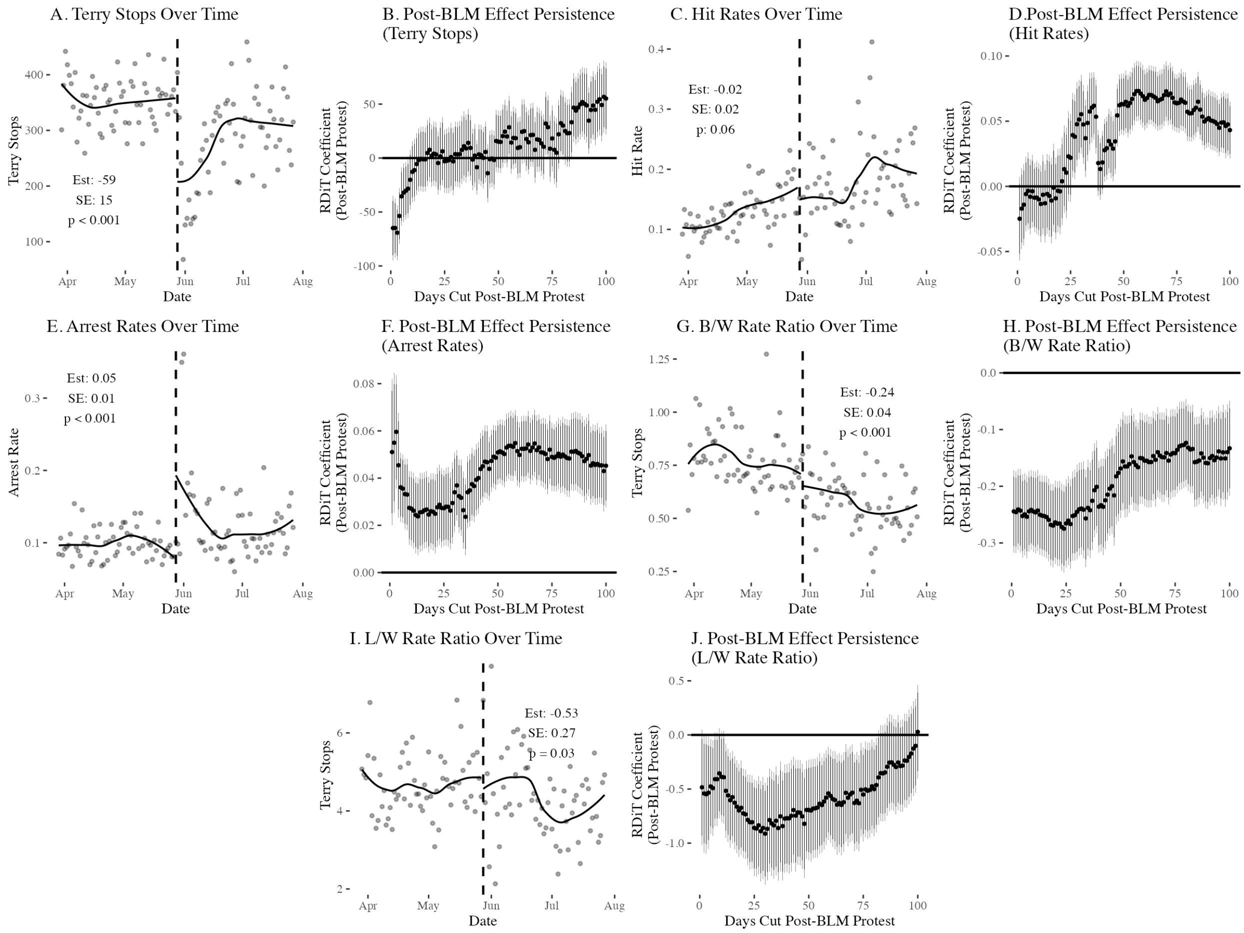

Figure I70: San Diego Replication. Panels A, C, E, and G characterize daily terry stops, hit rates, arrest rates, and reported violent crime from 911 calls (y-axis) over time (x-axis). Annotations for Panels A, C, E, and G are RDiT estimates characterizing the discontinuous post-*BLM protest* coefficient for the respective outcomes (linear polynomial, uniform kernel, mean-squared optimal bandwidth selection by Calonico, Cattaneo, and Titiunik (2015)). Panels B, D, F, and H characterize RDiT post-*BLM protest* coefficient after cutting 1-100 days post-*BLM protest* (but keeping days after intact). 95% CIs displayed, robust SEs reported.

Figure I70 characterizes the results from analyzing San Diego stop and 911 call data to test *Hypotheses 1-2*. To test *Hypothesis 1-2*, we collect San Diego Police Department (SDPD) terry stop data from San Diego's open data website.[30] With these data, we generate daily-

[30] `https://data.sandiego.gov/datasets/?department=police`

level measures of the count of terry stops, hit rates,[31] arrest rates, black/white rate ratios, and Latino/white rate ratios. The San Diego data are aggregated at the stop/person-level. This is because multiple people may be involved in a given stop (e.g. a police officer stopping a group of 3 people). Therefore, the daily time series of terry stops is aggregated from stop-level data whereas the the daily time series of hit rates, arrest rates, and the rate ratios are aggregated from person/stop-level data.

Despite being a Republican-controlled city at the time of the 2020 BLM protests, our results in San Diego are largely consistent with our main results derived from Democrat-controlled cities, with some minor exceptions.

First, consistent with *Hypothesis 1*, SDPD terry stops discontinuously decrease post-*BLM protest*. Descriptively, this is clear in Figure I70, Panel A. RDiT estimates also suggest the *BLM protest* discontinously decreased the number of terry stops by 59, 70% of the pre-*BLM protest* outcome standard deviation (see annotation on Figure I70, Panel A). However, unlike our main results, the *BLM protest* does not appear to persistently reduce SDPD terry stops. After cutting 1-100 days immediately post-*BLM protest* from the data (but leaving data from the days after intact), the decrease in terry stops reverts to the pre-*BLM protest* equilibrium within 15 days (Figure I70, Panel B).

Second, consistent with our main results, we find mixed evidence for *Hypothesis 2a* and *2b*. On the one hand, consistent with *Hypothesis 2b*, there is no discontinuous increase in the hit rate post-*BLM protest* (Figure I70, Panel C). However, consistent with *Hypothesis 2a*, there is a discontinuous increase in the arrest rate (Figure I70, Panel E). RDiT estimates suggest the *BLM protest* increased the arrest rate by 5 percentage points, equivalent to 150% of the pre-*BLM protest* outcome standard deviation. The increase in arrest rates post-*BLM protest* was persistent, suggesting the post-*BLM protest* coefficient is not driven by dynamics intrinsic to the behavior of protesters. The discontinuous increase in arrest rates post-*BLM protest* remained after cutting 1-100 days immediately post-*BLM protest* from the data Figure I70, Panel F).

Third, consistent with our main results, we find additional support for *Hypothesis 2a* analyzing racial disparities in SDPD terry stops. The Black/white and Latino/white stop rate ratios discontinuously decrease post-*BLM protest* (Figure I70, Panels G and I). RDiT estimates suggest the *BLM protest* discontinuously decreased the Black/white and Latino/white stop rate ratios by 0.24 and 0.53 respectively, equivalent to 180% and 71% of the pre-*BLM protest* outcome standard deviations. These decreases were persistent, again suggesting the post-*BLM protest* coefficient is not driven by dynamics intrinsic to the behavior of protestors. The discontinuous decrease in Black/white stop rate ratios remained after cutting 1-100 days immediately post-*BLM protest* from the data Figure I70, Panel H). Likewise, the discontinuous decrease in Latino/white stop rate ratios remained up to 80 days immediately post-*BLM protest* (Figure I70, Panel J).

In summary, like our main results, there is evidence consistent with *Hypothesis 1*, but mixed evidence consistent with *Hypothesis 2a* and *Hypothesis 2b*. On balance, there is a decrease in policing activity, but only briefly. Concomitantly, there is no discontinuous shift

[31] We use the SDPD definition of contraband to measure hit rates conditional on a terry stop. SDPD defines contraband as alcohol, ammunition, cellphones/electronic devices, drug paraphernalia, drugs/narcotics, firearms, money, stolen property, non-firearm weapons (see `https://www.sandiego.gov/sites/default/files/sdpd-ripa-presentation-220124.pdf`).

in hit rates, but there is a discontinuous (and persistent) increase in arrest rates in addition to a discontinuous (and persistent) decrease in racial disparities.