

Panic on the Streets of London: Police, Crime, and the July 2005 Terror Attacks[†]

By MIRKO DRACA, STEPHEN MACHIN, AND ROBERT WITT*

Terrorism is arguably the single most significant topic of political discussion of the past decade. In response, a small economics literature has begun to investigate the causes and impacts of terrorism (see Alan B. Krueger 2007 for a summary, or Krueger and Jitka Maleckova 2003 for some empirical work). Terror attacks, or the threat thereof, have also been considered in research on one important area of public policy, namely, the connections between crime and policing. Some recent studies (such as Rafael Di Tella and Ernesto Schargrodsky 2004; Jonathan Klick and Alexander Tabarrok 2005) have used terrorism-related events to look at the crime-police relationship, since terror attacks can induce an increased police presence in particular locations. This deployment of additional police can be used, under certain conditions, to test whether increased police activity reduces crime.

In this paper, we also consider the crime-police relationship before and after a terror attack, but in a very different context to other studies, by looking at the increased security presence following the terrorist bombs that hit central London in July 2005. Our application is a more general one than the other studies in that it covers a large metropolitan area following one of the most significant and widely known terror attacks of recent years. The scale of the security response in London after these attacks provides a potentially useful setting to examine the relationship between crime and police.

Moreover, and unlike the other studies in this area, we have very good data on police deployment. We can use these to identify the magnitude of the causal impact of police on crime.¹ A major strength of this paper is, therefore, that we are able to offer explicit instrumental variable-based estimates of the crime-police elasticity, which can be compared to other estimates like those of Steven D. Levitt (1997), Hope Corman and H. Naci Mocan (2000), and Di Tella and Schargrodsky's (2004) implied elasticity.² In fact, the sharp discontinuity in police deployment we are

* Draca: Centre for Economic Performance, London School of Economics, Houghton Street, London, WC2A 2AE, United Kingdom (e-mail: m.draca@lse.ac.uk); Machin: Department of Economics, University College London, Gower Street, London, WC1E 6BT, and Centre for Economic Performance, London School of Economics, Houghton Street, London, WC2A 2AE (e-mail: s.machin@ucl.ac.uk); Witt: Department of Economics, University of Surrey, Guildford, Surrey, GU2 7XH (e-mail: r.witt@surrey.ac.uk). We would like to thank the Metropolitan Police and Gerry Weston at Transport for London for assistance with the data used in this study. Helpful comments were received from three referees, Joshua Angrist, Christian Dustmann, Radha Iyengar, Alan Manning, Enrico Moretti, Ferdinand Rauch, from seminar participants at Aberdeen, Essex, LSE, UCL, and the Paris School of Economics, and from participants in the 2007 Royal Economic Society annual conference held in Warwick, the 2007 American Law and Economics Association annual meeting at Harvard, the 2007 European Economic Association annual conference in Budapest, and the NBER Inter American Seminar in Economics in Buenos Aires. Amy Challen and Richard Murphy contributed valuable research assistance.

[†] To view additional materials, visit the article page at <http://www.aeaweb.org/articles.php?doi=10.1257/aer.101.5.2157>.

¹ Neither Di Tella and Schargrodsky (2004) nor Klick and Tabarrok (2005) have data on police activity.

² While the Levitt (1997) paper is well known and widely cited, Justin McCrary's (2002) comment highlights some concerns about the data and the approach used (see also Levitt's 2002 response).

able to identify using this data means we are able to pin down this causal relation between crime and police very precisely. The natural experiment we consider also has some important external validity in the sense that it involves the deployment of a clear “deterrence technology” (that is, more police on the streets) rather than a measure of increased expenditures on police (e.g., as in William N. Evans and Emily G. Owens 2007; Machin and Olivier Marie forthcoming). Arguably, this type of visible increase in police deployment is the main type of policy mechanism under discussion in public debates about the funding and use of police resources.

Furthermore, the effectiveness of police is important in the context of a large criminological literature that has generally failed to find significant impacts of police on crime, even in quasi-experimental studies. Lawrence W. Sherman and David Weisburd (1995) review some of the conclusions from this work. Michael R. Gottfredson and Travis Hirschi (1990, p. 270) state “no evidence exists that augmentation of police forces or equipment, differential police strategies, or differential intensities of surveillance have an effect on crime rates.” Similarly emphatic arguments are made in Carl Klockars (1983).

The current paper focuses on what happened to criminal activity following a large and unanticipated increase in police presence. The scale of the change in police deployment that we study is much larger than in any of the other work in the crime-police research field. Indeed, results reported below show that police activity in central London increased by over 30 percent in the six weeks following the July 7 bombings as part of a police deployment policy stylishly titled “Operation Theseus” by the authorities. This police intervention represented the deployment of a very strong deterrence technology. The coverage of police was more sustained, widespread and complete than in the previous studies in the existing literature. We therefore view the scale of this change as important in addressing the paradox of the criminology literature discussed above where it proves hard to detect crime reductions linked to increased police presence. This is particularly the case since, during the time period when police presence was heightened, crime fell significantly in central London relative to outer London. Both the timing of the crime reductions and the types of crime that were more affected make us confident that this research approach identifies a causal impact of police on crime. Moreover, when police deployments returned to their pre-attack levels some six weeks later, the crime rate rapidly returned to its pre-attack level. Exploiting these sharp discontinuities in police deployment, we estimate an elasticity of crime with respect to police of approximately -0.3 to -0.4 , so that a 10 percent increase in police activity reduces crime by around 3 to 4 percent. Furthermore, we are unable to find evidence of either temporal or contemporaneous spatial displacement effects arising from the six-week police intervention.

A crucial part of identifying a causal impact in this type of setting is establishing the exclusion restriction that terrorist attacks affect crime through the post-attack increase in police deployment, and not via other observable and unobservable factors correlated with the attack or shock. Establishing this is important to generate credibility that our findings inform the crime-police debate rather than being just about an episode where a terror attack occurred. In this regard, the police deployment data we use are invaluable, as their availability makes it possible to distinguish the impact of police on crime from any general impact of the terrorist attack. In

particular, the research design features two discontinuities related to the police intervention. The first is the introduction of the geographically focused police deployment policy in the week of the terrorist attack. The immediate period surrounding the introduction of the policy was also characterized by a series of potentially correlated observable and unobservable shocks related to the attack. In contrast, the second discontinuity associated with the withdrawal of the policy occurred in a very different context. In this case, the observable and unobservable shocks associated with the attack were still in effect and dissipating gradually. Crucially, though, the police deployment was discretely “switched off” after a six-week period and we observe an increase in crime that is timed exactly with this change. Thus, we argue that it is difficult to attribute such a clear change in crime rates to observable and unobservable shocks arising from the terrorist attacks. If these types of shocks significantly affected crime rates, we would expect this to continue even as the police deployment was withdrawn. Indeed, an interesting feature of our empirical results is just how clearly and definitively crime seems to respond to a police presence.

The rest of the paper is organized as follows. Section I describes the events of July 2005 and goes over the main modeling and identification issues. In Section II we describe the data and provide an initial descriptive analysis. Section III presents the statistical results and a range of additional empirical tests. Section IV concludes.

I. Crime, Police, and the London Terror Attacks

A. *The Terror Attacks*

In July 2005 London’s public transport system was subject to two waves of terror attacks. The first occurred on Thursday, July 7, and involved the detonation of four bombs. The 32 boroughs of London are shown in Figure 1. Three of the bombs were detonated on London Underground (the tube) train carriages near the stations of Russell Square (in the borough of Camden), Liverpool Street (Tower Hamlets), and Edgware Road (Kensington and Chelsea). A fourth bomb was detonated on a bus in Tavistock Square, Bloomsbury (Camden). The second wave of attacks occurred two weeks later on July 21, consisting of four unsuccessful attempts at detonating bombs on trains near the underground stations of Shepherds Bush (Kensington and Chelsea), the Oval (Lambeth), Warren Street (Westminster), and on a bus in Bethnal Green (Tower Hamlets). Despite the failure of the bombs to explode, this second wave of attacks caused much turmoil in London. There was a large manhunt to find the four men who escaped after the unsuccessful attacks, and all of them were captured by July 29.

B. *Terror Attacks, Crime, and Correlated Shocks*

Di Tella and Schargrodsky (2004) were first to use police allocation policies in the wake of terror attacks to circumvent the endogeneity problem of crime and police. Using a July 1994 terrorist attack that targeted the main Jewish center in Buenos Aires, they show that motor vehicle thefts fell significantly in areas where extra police were subsequently deployed compared to areas several blocks away that did not receive



FIGURE 1. MAP OF LONDON BOROUGHS

Notes: This map represents the 32 boroughs of London. The treatment group for Operation Theseus police intervention includes: Camden, Kensington and Chelsea, Islington, Tower Hamlets, and Westminster. See Table A1 of the online Appendix for descriptive statistics on crime levels for the treatment and comparison groups.

Source: EDINA (Edinburgh National Academic Data Centre) archive.

extra protection. Their effect is large (approximately a 75 percent reduction in thefts relative to the comparison group) but also extremely local, with no evidence that the police presence reduced crime one or two blocks away from the protected areas. Another study by Klick and Tabarrok (2005) uses terror alert levels in Washington, DC, to make inferences about the crime-police relationship. The deployments they consider cover a more general area but (as already discussed) in the end are speculative since they are not able to quantify them with data on police numbers or hours.

Both papers touch on the issue of correlated shocks to observables and unobservables. However, in our case of London, this could be a greater concern since the terrorist attacks were a more significant, dislocating event for the city. Therefore, in thinking about the question of correlated shocks, it is helpful first to consider a basic equation, specified in levels, which describes the determinants of the crime rate in a set of geographical areas (in our case, London boroughs) over time:

$$(1) \quad C_{jt} = \delta P_{jt} + \lambda \mathbf{X}_{jt} + \mu_j + \tau_t + \varepsilon_{jt},$$

where C_{jt} is the crime rate for borough j in time period t , P_{jt} the level of police deployed, and \mathbf{X}_{jt} is a vector of control variables that could comprise observable or unobservable elements. The next set of terms is: μ_j , a borough level fixed effect; τ_t , a common time effect (for example, to capture common weather or economic shocks), and ε_{jt} , a random error term. In equation (1), t denotes weeks as we estimate weekly

crime equations, in which we are careful to recognise that crime displays a strong seasonal persistence.³

Consider a seasonally differenced version of equation (1), where the dependent variable is the change in the area crime rate relative to the rate in the same week of the previous year. This is highly important in crime modeling since crime is strongly persistent across areas over time. In practical terms, differencing eliminates the borough-level fixed effect, yielding

$$(2) \quad (C_{jt} - C_{j(t-k)}) = \delta(P_{jt} - P_{j(t-k)}) + \lambda(\mathbf{X}_{jt} - \mathbf{X}_{j(t-k)}) + (\tau_t - \tau_{t-k}) \\ + (\varepsilon_{jt} - \varepsilon_{j(t-k)}).$$

Note that the $\tau_t - \tau_{t-k}$ difference term can now be interpreted as the year-on-year change in factors that are common across all of the areas. By expressing this equation more concisely, we can make the correlated shocks issue explicit as follows:

$$(3) \quad \Delta_k C_{jt} = \delta \Delta_k P_{jt} + \lambda \Delta_k \mathbf{X}_{jt} + \Delta_k \tau_t + \Delta_k \varepsilon_{jt},$$

where Δ is a difference operator with k indexing the order of the seasonal differencing.

Using this framework we can carefully consider how a terrorist attack—which we can denote generally as Z —affects the determinants of crime across areas. Following the argument in the papers discussed above, the terror attack Z affects ΔP_{jt} , shifting police resources in a way that one can hypothesize is unrelated to crime levels. This hypothesis is, of course, a crucial aspect of identification that needs serious consideration. For example, it is possible that Z could affect the elements of ΔX_{jt} , creating additional channels via which terrorist attacks could influence crime rates.

What are these potential impacts or channels? The economics of terrorism literature stresses that the impacts of terrorism can be strong, but generally turn out to be temporary (Patrick Lenain, Marcos Bonturi, and Vincent Koen 2002; Nicholas Bloom 2009) in that economic activity tends to recover and normalize itself fairly rapidly. Of course, a sharp but temporary shock would still have ample scope to intervene in our identification strategy by affecting crime in a way that is correlated with the police response. In particular, three channels demand consideration. First, there is the physical dislocation caused by the attack. A number of tube stations were closed and many Londoners changed their mode of transport after the attacks (e.g., from the tube to buses or bicycles). This would have reshaped travel patterns and could have affected the potential supply of victims for criminals in some areas. Second, the volume of overall economic activity was affected. Studies on the aftermath of the attack indicate that both international and domestic tourism fell after the attacks, as measured by hotel vacancy rates, visitor spending data, and counts of domestic day trips (Greater London Authority 2005). Finally, there may be a psychological impact on individuals in terms of their attitudes

³ These types of effects could prevail where seasonal patterns affect different boroughs with varying levels of intensity. For example, the central London boroughs are more exposed to fluctuations due to tourism activity, and exhibit sharper seasonal patterns with respect to crime.

toward risk. As Gary Becker and Yona Rubinstein (2004) outline, this influences observable travel decisions as well as more subtle unobservable behavior.

To summarize, we think of these effects as being manifested in three elements of the \mathbf{X}_{jt} vector outlined above:

$$(4) \quad \mathbf{X}_{jt} = [\mathbf{X}_{jt}^1, \mathbf{X}_{jt}^2, \theta_{jt}].$$

In (4), \mathbf{X}_{jt}^1 is a set of exogenous control variables (observable to researchers) that includes observable factors such as area-level labor market conditions that change slowly and are unlikely to be immediately affected by terrorist attacks (if at all). The second \mathbf{X}_{jt}^2 vector represents the observable factors that change more quickly and are therefore vulnerable to the dislocation caused by terrorist attacks. As discussed above, here we are thinking primarily of factors such as travel patterns that could influence the potential supply of victims to crime across areas. The final element θ_{jt} , then, captures an analogous set of unobservable factors that are susceptible to change due to the terrorist attack. In the spirit of Becker and Rubinstein's (2004) discussion, the main factor to consider here is fear, or how individuals handle the risks associated with terrorism. For example, it is plausible that, in the wake of the attacks, commuters in London became more vigilant to suspicious activity in the transport system and in public spaces. This vigilance would have been focused mainly on potential terrorist activity, but one might expect that this type of cautious behavior could have a spillover onto crime.

The implications of these correlated shocks for our identification strategy can now be clearly delineated. For our exclusion restriction to hold, it needs to be shown that the terrorist attack Z affected the police deployment in a way that can be separately identified from Z 's effect on other observable and unobservable factors that can influence crime rates. Practically, we show this later in the paper by mapping the timing and location of the police deployment shock and comparing it to the profiles of the competing observable and unobservable shocks.

C. Possible Displacement Effects

Another issue that could affect identification is crime displacement. Since the police intervention affected the costs of crime across locations and time, it may be that criminals take these changes into account and adjust their behavior. This raises the possibility that criminal activity was either diverted into other areas (e.g., the comparison group of boroughs) during Operation Theseus or postponed until after the extra police presence was withdrawn. The implication is that simple difference-in-differences (DiD) estimates of the police effect on crime would be upwardly biased if these offsetting spatial displacement effects were not taken into account. Temporal displacement can have the opposite effect, and we discuss this more in the final empirical section.

II. Data Description and Initial Descriptive Analysis

A. Data

We use daily police reports of crime from the London Metropolitan Police Service (LMPS) before and after the July 2005 terrorist attacks. Our crime data cover the

period from January 1, 2004, to December 31, 2005, and are aggregated up from ward to borough level and from days to weeks over the two-year period. There are 32 London boroughs as shown on the map in Figure 1.⁴ There are also monthly borough-level data available over a longer time period that we use for some robustness checks.

The basic street-level policing of London is carried out by 33 Borough Operational Command Units (BOCUs), which operate to the same boundaries as the 32 London borough councils, apart from one BOCU, which is dedicated to Heathrow Airport. We have been able to put together a weekly panel covering 32 London boroughs over two years, giving 3,328 observations. Crime rates are calculated on the basis of population estimates at borough level, supplied by the Office of National Statistics (ONS) online database.⁵

The police deployment data are at borough level and were produced under special confidential data-sharing agreements with the LMPS. The main data source used is CARM (Computer Aided Resource Management), the police service's human resource management system. This records hours worked by individual officers on a daily basis. We aggregate the deployment data to borough level since the CARM data are mainly defined at this level. However, there is also useful information on the allocation of hours worked by incident and/or police operation.⁶ While hours worked are available according to officer rank, our main hours measure is based on total hours worked by all officers in the borough adjusted for this reallocation effect. In addition to crime and deployment, we have obtained weekly data on tube journeys for all stations from Transport for London (TFL). It is daily borough-level data aggregated up to weeks based on entries into and exits from tube stations. Finally, we also use data from the UK Labour Force Survey (LFS) to provide information on local labor market conditions.

B. Initial Approach

Our analysis begins by looking at what happened to police deployment and crime before and after the July 2005 terror attacks in London using a difference-in-differences (DiD) approach. This rests upon defining a treatment group of boroughs in central and inner London where the extra police deployment occurred, and comparing their crime outcomes to the other, nontreated boroughs. The police hours data we use facilitate the development of this approach, with two features standing out. First, the data allow us to measure the increase in total hours worked in the period after the attacks. The increase in total hours was accomplished through the increased use of overtime shifts across the police service, and this policy lasted approximately six weeks. Second, the police data contain a special resource allocation code denoted as Central Aid. This code allows us to identify how police hours worked were geographically reallocated over the six-week period. For example, we

⁴ The City of London has its own police force and so this small area is excluded from our analysis.

⁵ Online Appendix Table A1 shows some summary statistics on the crime data.

⁶ Since the CARM information is also used for calculating police pay, it is considered a very reliable measure of police activity. We gained access to this data after repeated inquiries to the LMPS. The main condition for access was that we not reveal any strategic information about ongoing or individual, borough-specific police deployment policies.

can identify how hours worked by officers stationed in the outer London boroughs were reallocated to public security duties in central and inner London. The extra hours were mainly reallocated to the boroughs of Westminster, Camden, Islington, Kensington and Chelsea, and Tower Hamlets, with individual borough allocations proportional to the number of tube stations in the borough.⁷ These boroughs either contained the sites of the attacks or featured many potential terrorist targets, such as transport nodes or significant public spaces. Using these two features of the data, we are able to define a treatment group comprising the five named boroughs. A map showing the treatment group is given in Figure 1. In most of the descriptive statistics and modeling below, we use all other boroughs as the comparison group in order to simplify the analysis.

What did the extra police deployment in the treated boroughs entail? The number of mobile police patrols were greatly increased and officers were posted to guard major public spaces and transport nodes, particularly tube stations. In areas of central London where many stations were located, this resulted in a highly visible police presence, and this is confirmed by public surveys conducted at the time.⁸ This high visibility potentially exerted a deterrent effect on public, street-level crimes such as thefts and violent assault. We test for this prediction in the empirical work.

C. Basic Difference-in-Differences

In Table 1 we compare what happened to police deployment and to total crime rates before and after the July 2005 terror attacks in the treatment group boroughs, as compared to all other boroughs. Police deployment is measured in a similar way to crime rates; that is, we normalize police hours worked by the borough population. Following the discussion in Section II we define the before and after periods in year-on-year, seasonally adjusted terms. This ensures that we are comparing like-with-like in terms of the seasonal effects prevailing at a given time of the year. For example, looking at Table 1, the crime rate of 4.03 per 1,000 population in panel B represents the treatment group crime rate in the period from July 8, 2004, until August 19, 2004. The post-period or “policy on” period then runs from July 7, 2005, until August 18, 2005, with a crime rate of 3.59.⁹ Thus, by taking the difference between these “pre” and “post” crime rates we are able to derive the year-on-year, seasonally adjusted change in crime rates and police hours. These are then differenced across the treatment ($T = 1$) and comparison ($T = 0$) groups to get the customary DiD estimate.

⁷ We say “mainly reallocated” due to the fact that some mobile patrols crossed into adjacent boroughs, and because some bordering areas of boroughs were the site of some small deployments. A good case here is the southern tip of Hackney borough (between Islington and Tower Hamlets). However, the majority of Hackney was not treated by the policy (since this borough is notoriously lacking in Tube station links), so we exclude it from the treatment group.

⁸ Table A2 of the online Appendix reports the results of a survey of London residents in the aftermath of the attacks. Approximately 70 percent of respondents from inner London attested to a higher police presence in the period since the attacks. The lower percentage reported by outer London residents also supports the hypothesis of differential deployment across areas.

⁹ The one-day difference in calendar date across years ensures we compare the same days of the week.

TABLE 1—POLICE DEPLOYMENT AND MAJOR CRIMES, DIFFERENCE-IN-DIFFERENCES, 2004–2005

	Panel A. Police deployment (Hours worked per 1,000 population)			Panel B. Crime rate (Crimes per 1,000 population)		
	Pre-period	Post-period	Difference (post–pre)	Pre-period	Post-period	Difference (post–pre)
	(1)	(2)	(3)	(4)	(5)	(6)
$T = 1$	169.46	242.29	72.83	4.03	3.59	–0.44
$T = 0$	82.77	84.95	2.18	1.99	1.97	–0.02
Difference-in-differences (levels)			70.65*** (7.52)			–0.43** (0.16)
Difference-in-differences (logs)			0.35*** (0.04)			–0.11*** (0.04)

Notes: Post-period defined as the six weeks following 7/7/2005. Pre-period defined as the six weeks following 8/7/2004. Weeks defined in a Thursday–Wednesday interval throughout to ensure a clean pre and post split in the 2005 attack weeks. Treatment group ($T = 1$) defined as boroughs of Westminster, Camden, Islington, Tower Hamlets, and Kensington–Chelsea. Comparison group ($T = 0$) defined as other boroughs of London. Police deployment defined as total weekly hours worked by police staff at borough level. Standard errors are in parentheses.

*** Significant at the 1 percent level.

** Significant at the 5 percent level.

* Significant at the 10 percent level.

The first panel of Table 1 shows unconditional DiD estimates for police hours. It is clear that the treatment boroughs experienced a very large relative change in police deployment. Per capita hours worked increased by 34.6 percent in the DiD (final row, column 3). Arguably, the *composition* of this relative change is almost as important for our experiment as the scale. The relative change was driven by an increase in the treatment group (of 72.8 hours per capita) with little change in hours worked for the comparison group (only 2.2 hours per capita more). This was feasible because of the large number of overtime shifts worked. In practice, this means that, while there was a diversion of police resources from the comparison boroughs to the treatment boroughs, the former areas were able to keep their levels of police hours constant. Obviously, this *ceteris paribus* feature greatly simplifies our later analysis of displacement effects, since we do not have to deal with the implications of a zero-sum shift of resources across areas. The next panel of Table 1 deals with the crime rates. It shows that crime rates fell by 11.1 percent in the DiD (final row, column 6). Again, this change is driven by a fall in treatment group crime rates and a steady crime rate in the comparison group. This is encouraging since it is what would be expected from the type of shift we have just seen in police deployment.

Weekly police deployment and crime rates are shown in panels A and B of Figure 2. Here we do two things. First, we normalize crime rates and police hours across the treatment and comparison groups by their level in week one of our sample (i.e., January 2004). This rescales the levels in both groups so that we can compare their evolution over time. Second, we mark out the attack (“policy-on”) period in 2005, along with the comparison period in the previous year. As panel A shows, this reveals a clear, sharp discontinuity in police deployment. Police hours in the treatment group rise immediately after the attack and fall sharply at the end of the six weeks of Operation Theseus.

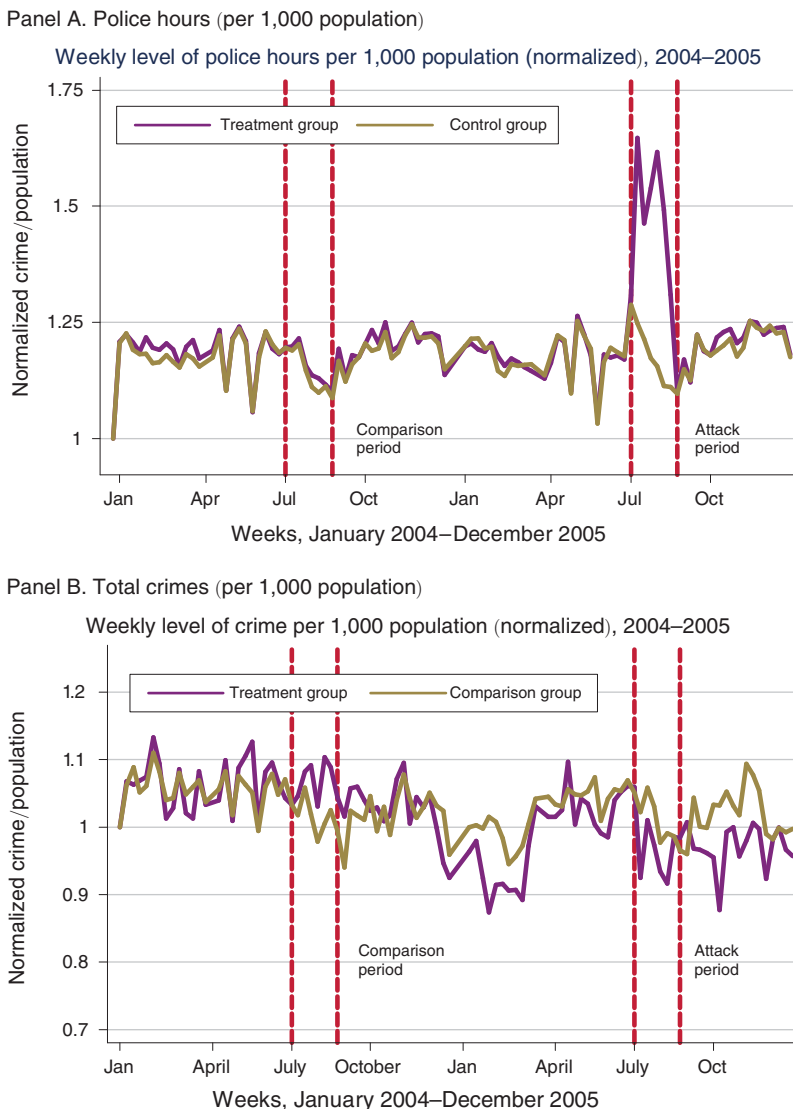


FIGURE 2. POLICE HOURS AND TOTAL CRIME (LEVELS) 2004–2005, TREATMENT VERSUS COMPARISON GROUP

Notes: This figure plots levels of police and crime for the treatment and comparison groups. Horizontal axis covers the period from January 2004 to January 2006. The values of police and crime have been normalized relative to the values in the first week of January 2004. Treatment and comparison groups defined as per Figure 1.

The visual evidence for the crime rate in panel B of Figure 2 is less decisive because the weekly crime rates are clearly more volatile than the police hours data. This is to be expected insofar as police hours are largely determined centrally by policymakers, while crime rates are essentially the outcomes of decentralized activity. This volatility does raise the possibility that the fall in crime rates seen in the Table 1 DiD estimates may simply be due to naturally occurring, short-run time series volatility rather than the result of a policy intervention—a classic problem in the literature (John J. Donohue 1998). After the correlated shocks issue, this is

probably the biggest modeling issue in the paper and we deal with it extensively in the next section.

III. Statistical Models of Crime and Police

In this section we present our statistical estimates. We begin with a basic set of estimates and then move on to focus on specific issues to do with different crime types, timing, correlated shocks and displacement effects.

A. Statistical Approach

The starting point for the statistical work is a DiD model of crime determination. We have borough-level weekly data for the two calendar years 2004 and 2005. The terror attack variable (Z as discussed above) is specified as an interaction term $T_b \times POST_t$, where T denotes the treatment boroughs and $POST$ is a dummy variable equal to one in the post-attack period.

In this setting the basic reduced-form seasonally differenced weekly models for police deployment and crime (with lower case letters denoting logs) are

$$(5) \quad \Delta_{52}p_{bt} = \alpha_1 + \beta_1 POST_t + \delta_1(T_b \times POST_t) + \lambda_1 \Delta_{52}x_{bt} + \Delta_{52}\varepsilon_{1bt},$$

$$(6) \quad \Delta_{52}c_{bt} = \alpha_2 + \beta_2 POST_t + \delta_2(T_b \times POST_t) + \lambda_2 \Delta_{52}x_{bt} + \Delta_{52}\varepsilon_{2bt}.$$

Due to the highly seasonal nature of crime noted above, the equations are differenced across weeks of the year (hence the $k = 52$ subscript in the Δ_k differences). The key parameters of interest are δ_1 and δ_2 , the seasonally adjusted difference-in-differences estimates of the impact of the terror attacks on police deployment and crime.

These reduced-form equations can be combined to form a structural model relating crime to police deployment, from which we can identify the causal impact of police on crime. The structural equation is

$$(7) \quad \Delta_{52}c_{bt} = \alpha_3 + \beta_3 POST_t + \delta_3 \Delta_{52}p_{bt} + \lambda_3 \Delta_{52}x_{bt} + \Delta_{52}\varepsilon_{3bt},$$

where the variation in police deployment induced by the terror attacks identifies the causal impact of police on crime. The first-stage regression is equation (5) above, and so equation (7) is estimated by instrumental variables (IV) where the $T_b \times POST_t$ variable is used as the instrument for the change in police deployment. Here, the structural parameter of interest, δ_3 (the coefficient on police deployment), is equal to the ratio of the two reduced-form coefficients, so that $\delta_3 = \delta_2/\delta_1$.

Finally, note that in some of the reduced-form specifications that we consider below, we split the $POST_t \times T_b$ into two distinct post-7/7 time periods so as to distinguish the “post-policy” period after the end of Operation Theseus. This term is added in order to test directly for any persistent effect of the police deployment, and importantly to focus explicitly on the second “experiment” when police levels fell sharply back to their pre-attack levels. Thus, the reduced forms in (5) and (6) now become

$$(8) \quad \Delta_{52}p_{bt} = \alpha_4 + \beta_4 POST_t + \delta_{41}(T_b \times POST_t^1) \\ + \delta_{42}(T_b \times POST_t^2) + \lambda_4 \Delta_{52}x_{bt} + \Delta_{52}\varepsilon_{4bt},$$

$$(9) \quad \Delta_{52}c_{bt} = \alpha_5 + \beta_5 POST_t + \delta_{51}(T_b \times POST_t^1) \\ + \delta_{52}(T_b \times POST_t^2) + \lambda_5 \Delta_{52}x_{bt} + \Delta_{52}\varepsilon_{5bt}.$$

In these specifications, $POST_t^1$ represents the six-week policy period immediately after the July 7 attack when the police deployment was in operation, while $POST_t^2$ covers the time period subsequent to the deployment until the end of the year (that is, from August 19, 2005, until December 31, 2005).¹⁰ Also, note that a test of $\delta_{41} = \delta_{42}$ (in the police equation, (8)) or $\delta_{51} = \delta_{52}$ (in the crime equation, (9)) amounts to a test of temporal variations in the initial six-week period directly after July 7 as compared to the remainder of the year.

B. Basic Difference-in-Differences Estimates

Table 2 provides the basic reduced forms for police (in panel A), for crime (in panel B) and OLS and structural IV results (in panel C) for the models outlined in equations (5)–(9). In the reduced forms, we specify three terms to uncover the DiD estimates of interest. Specifically, in column 1 of panels A and B, we include an interaction term that uses the full period from July 7, 2005, to December 31, 2005, to measure the post-attack period (in the table denoting $T_b \times POST_t$ from equations (5) and (6) as $T \times Post-Attack$). The adjacent columns (2–4) then split this period in two with one interaction term for the six-week Operation Theseus period (denoting $T_b \times POST_t^1$ from equations (8) and (9) as $T \times Post-Attack1$) and another for the remaining part of the year (denoting $T_b \times POST_t^2$ as $T \times Post-Attack2$). As already noted, the second term is useful for testing whether there were any persistent effects of the police deployment or any longer-term trends in the treatment group after police deployment fell back to its pre-attack levels.

The findings from the unconditional DiD estimates reported earlier are confirmed in the basic models in Table 2. The estimated coefficient on $T \times Post-Attack1$ in the reduced-form police equation shows a 34.1 percent increase in police deployment during Operation Theseus, and there is no evidence that this persists for the rest of the year (i.e., the $T \times Post-Attack2$ coefficient is statistically indistinguishable from zero). For the crime rate reduced form there is an 11.1 percent fall during the six-week policy-on period with minimal evidence of either persistence or a treatment group trend in the estimates for the $T \times Post-Attack2$ variable.¹¹ Despite this we include a full set of 32 borough-specific trends in the

¹⁰ As we discuss later, police deployment levels in London boroughs were returned to their pre-attack baselines at the end of Operation Theseus.

¹¹ While we have seasonally differenced the data, one may have concerns about possible contamination from further serial correlation. We follow Marianne Bertrand, Esther Dufló, and Sendhil Mullainathan (2004) and collapse the data before and after the attacks and obtain extremely similar results: the estimate (standard error) based on collapsed data comparable to the $T \times Post-Attack1$ estimate in panel B, column 3 of Table 2 was -0.112 (0.027).

TABLE 2—DIFFERENCE-IN-DIFFERENCES REGRESSION ESTIMATES, POLICE DEPLOYMENT AND TOTAL CRIMES, 2004–2005.

	Full (1)	Split (2)	+Controls (3)	+Trends (4)	
<i>Panel A. Police deployment (Hours worked per 1,000 population)</i>					
$T \times \text{Post-Attack}$	0.081*** (0.010)				
$T \times \text{Post-Attack1}$		0.341*** (0.028)	0.342*** (0.029)	0.356*** (0.027)	
$T \times \text{Post-Attack2}$		-0.001 (0.011)	0.001 (0.010)	0.014 (0.016)	
Controls	No	No	Yes	Yes	
Trends	No	No	No	Yes	
Number of boroughs	32	32	32	32	
Observations	1,664	1,664	1,664	1,664	
	Full (1)	Split (2)	+Controls (3)	+Trends (4)	
<i>Panel B. Total crimes (Crimes per 1,000 population)</i>					
$T \times \text{Post-Attack}$	-0.052** (0.021)				
$T \times \text{Post-Attack1}$		-0.111*** (0.027)	-0.109*** (0.027)	-0.056* (0.030)	
$T \times \text{Post-Attack2}$		-0.033 (0.027)	-0.031 (0.028)	0.024 (0.054)	
Controls	No	No	Yes	Yes	
Trends	No	No	No	Yes	
Number of boroughs	32	32	32	32	
Observations	1,664	1,664	1,664	1,664	
	<i>OLS Estimates</i>		<i>IV Estimates</i>		
	Levels (1)	Differences (2)	Full (3)	Split (4)	+Trends (5)
<i>Panel C. Structural form</i>					
ln(police hours)	0.785*** (0.053)				
$\Delta \ln(\text{police hours})$		-0.031 (0.051)	-0.641** (0.301)	-0.318*** (0.093)	-0.183*** (0.066)
Controls	Yes	Yes	Yes	Yes	Yes
Trends	No	No	No	No	Yes
Number of boroughs	32	32	32	32	32
Observations	3,328	1,664	1,664	1,664	1,664

Notes: All specifications include week fixed effects. Standard errors clustered by borough in parentheses. Weighted by borough population. “Full” post-period for baseline models in column 1 of panels A and B defined as all weeks after 7/7/2005 until 31/12/2005 attack inclusive. Weeks defined in a Thursday–Wednesday interval throughout to ensure a clean pre- and post-split in the attack weeks. $T \times \text{Post-Attack}$ is then defined as interaction of treatment group with a dummy variable for the post-period. $T \times \text{Post-Attack1}$ is defined as interaction of treatment group with a deployment “policy” dummy for weeks 1–6 following the July 7, 2005, attack. $T \times \text{Post-Attack2}$ is defined as treatment group interaction for all weeks subsequent to the main Operation Theseus deployment. Treatment group defined as boroughs of Westminster, Camden, Islington, Tower Hamlets, and Kensington-Chelsea. Police deployment defined as total weekly hours worked by all police staff at borough level. Controls based on Quarterly Labour Force Survey (QLFS) data and include: borough unemployment rate, employment rate, males under 25 as proportion of population, and whites as proportion of population (following QLFS ethnic definitions).

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

specifications in column 4 to test robustness. The crime rate coefficient for the Operation Theseus period is halved, but the interaction term is still significant, indicating that there was a fall in crime during this period that was over and above that of any combination of trends.

Estimates of the structural model are given in panel C of Table 2. The OLS estimates are reported in columns 1 and 2. In column 1—labelled “levels”—a pooled cross-sectional regression shows a strongly significant positive coefficient on the police deployment variable, clearly showing the well-known problems of reverse causation. In column 2 we show estimates from a seasonally differenced version of this OLS regression, reporting a negligible, insignificant coefficient. This reflects the limited year-on-year change in police hours to be found when the seasonal difference is taken.

However, it is clear from the results reported in panels A and B of Table 2 that the coincident nature of the respective timings of the increase in police deployment and the fall in crime suggests that increased security presence lowered crime. The final three columns of panel C of the table therefore show estimates of the causal impact of increased deployment on crime. Column 3 shows the basic IV estimate where the post-attack effects are constrained to be time-invariant. Columns 4 and 5 allow for time variation to identify a more local causal impact. The instrumental variable estimates are precisely determined owing to the strength of the first-stage police regressions in panel A of the table. The preferred estimate with time-varying terror attack effects (reported in column 4) shows an elasticity of crime with respect to police of around -0.32 . This implies that a 10 percent increase in police activity reduces crime by around 3.2 percent. The magnitude of these causal estimates is similar to the small number of causal estimates found in the literature (they are also estimated much more precisely in statistical terms because of the very sharp discontinuity in police deployment that occurred). Levitt's (1997) study found elasticities in the -0.43 to -0.50 range, while Corman and Mocan (2000) estimated an average elasticity of -0.45 across different types of offences, and Di Tella and Schargrodsy (2004) reported an elasticity of motor vehicle thefts with respect to police of -0.33 .

C. Different Crime Types

Thus far the results use a measure of total crimes. Heterogeneity of the overall effect by the type of crime is potentially important, however. The pattern of the impact by crime type is an important falsification exercise. The main feature of Operation Theseus was a highly visible public deployment of police officers in the form of foot and mobile patrols, particularly around major transport hubs. We could therefore expect any police effect to be operating mainly through a deterrence technology based on greater visibility, generating an increase in the probability of detection for crimes committed in or around public places. As a result, the crime effect documented in Tables 1 and 2 should be concentrated in crime types more susceptible to this type of technology.

We therefore estimated reduced-form treatment effects across the six major crime categories defined by the Metropolitan Police—thefts, violent crimes, sexual offences, robbery, burglary, and criminal damage. Separate estimates by crime type

TABLE 3—TREATMENT EFFECTS BY MAJOR CRIME CATEGORY

<i>Panel A. Thefts, violence, and sex crimes</i>						
Crime category	Thefts		Violence		Sex crimes	
	(1)	(2)	(3)	(4)	(5)	(6)
$T \times \text{Post-Attack1}$	-0.139*** (0.044)	-0.082* (0.045)	-0.124*** (0.043)	-0.108*** (0.034)	-0.078 (0.123)	-0.102 (0.138)
$T \times \text{Post-Attack2}$	-0.017 (0.039)	0.044 (0.085)	-0.054 (0.032)	-0.038 (0.056)	-0.080 (0.082)	-0.094 (0.084)
Trends	No	Yes	No	Yes	No	Yes
Number of boroughs	32	32	32	32	32	32
Observations	1,664	1,664	1,664	1,664	1,664	1,664
<i>Panel B. Robbery, burglary, and criminal damage</i>						
Crime category	Robbery		Burglary		Criminal damage	
	(1)	(2)	(3)	(4)	(5)	(6)
$T \times \text{Post-Attack1}$	-0.131 (0.119)	-0.012 (0.129)	-0.035 (0.057)	-0.029 (0.067)	-0.047 (0.052)	-0.005 (0.041)
$T \times \text{Post-Attack2}$	-0.089 (0.098)	0.024 (0.149)	-0.093 (0.059)	-0.078 (0.075)	-0.018 (0.043)	0.020 (0.057)
Trends	No	Yes	No	Yes	No	Yes
Number of boroughs	32	32	32	32	32	32
Observations	1,664	1,664	1,664	1,664	1,664	1,664

Notes: All specifications include week fixed effects. Standard errors clustered by borough in parentheses. Weighted by borough population. $T \times \text{Post-Attack1}$ and $T \times \text{Attack2}$ defined as per Table 2. Treatment group also defined as per Table 2. See Table A6 in the online Appendix for definitions of the major crime categories in terms of the constituent minor crimes. Crime categories used follow the definitions provided by the MPS.

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

are reported in Table 3, which shows there to be important differences across these groups. The estimates show strongly significant effects for thefts and violent crimes. These are comprised of crimes such as street-level thefts (picking pockets, snatches, thefts from stores, motor vehicle-related theft, and tampering) as well as street-level violence (common assault, harassment, and aggravated bodily harm). Of considerable note is the lack of any effect for burglary. As a crime that occurs primarily at night and in private dwellings, this is arguably the crime category that is least susceptible to a public deterrence technology.

In Table 4 we aggregate these major categories into a group of crimes potentially susceptible to Operation Theseus (thefts, violent crimes, and robberies) and a group of remaining nonsusceptible crimes (burglary, criminal damage, and sexual offences). The susceptible crime estimates are given in part I of the table, and the nonsusceptible crime results in part II. The point estimate for our preferred susceptible crimes estimate is -0.132 (column 3, panel A), which compares to an estimate of -0.109 for total crimes in column 3, panel B, of Table 2. There is a much smaller (in absolute terms), statistically insignificant estimate of -0.023 for nonsusceptible crimes (reported in the column 3 of panel A in part II of Table 4). We therefore use this susceptible crimes classification as the main outcome variable in the remainder of our analysis. The estimated elasticity of susceptible crimes with respect to police deployment in the column 4, panel C, part I model of the table is -0.38 and is again very precisely determined.

TABLE 4, PART I—SUSCEPTIBLE CRIMES VERSUS NONSUSCEPTIBLE CRIMES, 2004–2005

	Reduced form				
	Full (1)	Split (2)	+Controls (3)	+Trends (4)	
<i>Panel A. I. Susceptible crimes</i>					
$T \times \text{Post-Attack}$	-0.056** (0.023)				
$T \times \text{Post-Attack1}$		-0.131*** (0.031)	-0.132*** (0.031)	-0.067* (0.035)	
$T \times \text{Post-Attack2}$		-0.033 (0.030)	-0.033 (0.030)	0.033 (0.063)	
Controls	No	No	Yes	Yes	
Trends	No	No	No	Yes	
Number of boroughs	32	32	32	32	
Observations	1,664	1,664	1,664	1,664	
	Structural form				
	Levels (1)	Differences (2)	Full (3)	Split (4)	+Trends (5)
<i>Panel B. OLS</i>			<i>Panel C. IV estimates</i>		
$\ln(\text{police hours})$	0.952*** (0.056)				
$\Delta \ln(\text{police hours})$		-0.019 (0.063)	-0.694** (0.336)	-0.383*** (0.105)	-0.223*** (0.074)
Controls	Yes	Yes	Yes	Yes	Yes
Trends	No	No	No	No	Yes
Number of boroughs	32	32	32	32	32
Observations	3,328	1,664	1,664	1,664	1,664

Notes: All specifications include week fixed effects. Standard errors clustered by borough in parentheses. Weighted by borough population. Susceptible crimes defined as: violence against the person; theft and handling; robbery. Nonsusceptible crimes defined as: burglary and criminal damage; sexual offenses. Treatment group definitions and $T \times \text{Post-Attack}$ terms defined as per Table 2. Controls also defined as per Table 2.

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

D. Timing

The previous section cited the volatility of the crime rates and timing in general as important issues. Given that we are using weekly data, there is a need to investigate to what extent short-term variations could be driving the results for our policy intervention. To test this, we take the extreme approach of testing every week for hypothetical or “placebo” policy effects. Specifically, we estimate the reduced-form models outlined in equations (5) and (6), defining a single week-treatment group interaction term for each of the 52 weeks in our data. We then run 52 regressions, each featuring a different week $\times T_b$ interaction, and plot the estimated coefficient and confidence intervals. The major advantage of this is that it extracts all the variation and volatility from the data in a way that reveals the implications for our main DiD estimates. Practically, this exercise is therefore able to test whether our six-week Operation Theseus effect is merely a product of time-series volatility or variation that is equally likely to occur in other subperiods.

We plot the coefficients and confidence intervals for all 52 weeks in panels A and B of Figure 3. Panel A shows the results for police hours repeating the

TABLE 4, PART II—SUSCEPTIBLE CRIMES VERSUS NONSUSCEPTIBLE CRIMES, 2004–2005

	Reduced form				
	Full (1)	Split (2)	+Controls (3)	+Trends (4)	
<i>Panel A. II. Nonsusceptible crimes</i>					
$T \times \text{Post-Attack}$	-0.048* (0.024)				
$T \times \text{Post-Attack1}$		-0.033 (0.026)	-0.023 (0.027)	-0.015 (0.031)	
$T \times \text{Post-Attack2}$		-0.053 (0.034)	-0.043 (0.037)	-0.033 (0.045)	
Controls	No	No	Yes	Yes	
Trends	No	No	No	Yes	
Number of boroughs	32	32	32	32	
Observations	1,664	1,664	1,664	1,664	
	Structural form				
	Levels (1)	Differences (2)	Full (3)	Split (4)	+Trends (5)
<i>Panel B. OLS</i>			<i>Panel C. IV estimates</i>		
$\ln(\text{police hours})$	0.327*** (0.046)				
$\Delta \ln(\text{police hours})$		-0.056 (0.094)	-0.597* (0.337)	-0.065 (0.078)	-0.012 (0.092)
Controls	Yes	Yes	Yes	Yes	Yes
Trends	No	No	No	No	Yes
Number of boroughs	32	32	32	32	32
Observations	3,328	1,664	1,664	1,664	1,664

Notes: All specifications include include week fixed effects. Standard errors clustered by borough in parentheses. Weighted by borough population. Susceptible crimes defined as: violence against the person; theft and handling; robbery. Nonsusceptible crimes defined as: burglary and criminal damage; sexual offenses. Treatment group definitions and $T \times \text{Post-Attack}$ terms defined as per Table 2. Controls also defined as per Table 2.

***Significant at the 1 percent level.


**Significant at the 5 percent level.

*Significant at the 10 percent level.

clear pattern seen in panel A of Figure 2 of the police deployment policy being switched on and off. (Note that precisely estimated treatment effects in this graph are characterized by confidence intervals that do not overlap the zero line.) The analogous result for the susceptible crime rate is then shown in panel B. The decreases in crime are less dramatic than the increases in police hours, but clearly the two closely coincide in timing. It is interesting to note that the pattern of six consecutive weeks of significant, negative treatment effects in the crime rate is not repeated in any other period of the data *except* Operation Theseus. This is impressive, as it shows that the effect of the policy intervention can be seen despite the noise and volatility of the weekly data.¹²

¹² As a further check on the issue of volatility, we made use of monthly, borough-level crime data available from 2001 onward (as the daily crime data we use to construct our weekly panel are available only since the beginning of 2004). These data allow us to examine whether there is a regular pattern of negative effects in the middle part of the year. In this exercise, we define year-on-year differences in susceptible crime for the July–August period over the range of intervals: 2001–2002, 2002–2003, 2003–2004, and 2004–2005. The results are shown in online Appendix Table A3. We find that a significant treatment effect in susceptible crimes is evident only for the 2004–2005 time period. This gives us further confidence that our estimate for this year is a unique event that cannot be likened to arbitrary fluctuations of previous years.

Panel A. Year-on-year change in police hours (per 1,000 population)



Panel B. Year-on-year change in susceptible crime rate

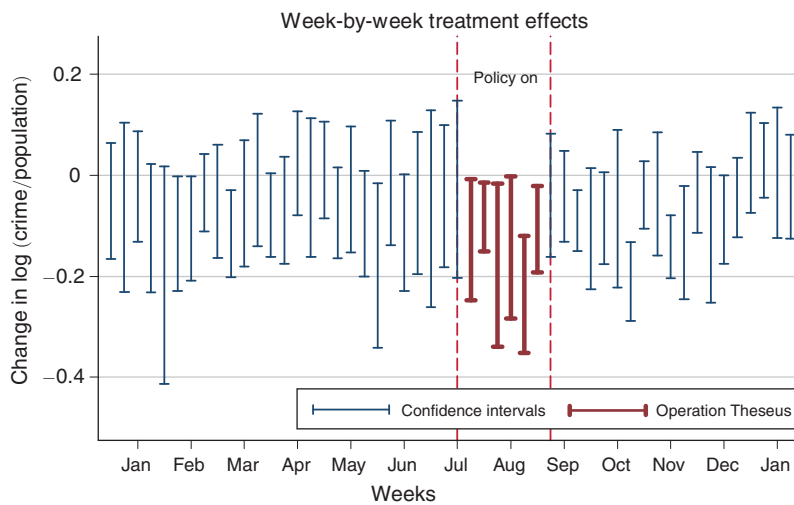


FIGURE 3. WEEK-BY-WEEK PLACEBO POLICY EFFECTS: POLICE HOURS AND SUSCEPTIBLE CRIMES

Notes: This figure plots the coefficients and confidence intervals for week-by-week treatment*week interactions from January 2005 to January 2006. These are estimated following the reduced-form specifications in the main body of the paper. Standard errors clustered by borough. Note that since this is year-on-year, seasonally differenced data, it reflects an underlying sample extending from January 2004 to January 2006.

E. Correlated Shocks

The discussion of timing has a direct bearing on the issue of correlated shocks outlined in Section II. In particular, it is important to examine the extent to which any shifts in correlated observables do or do not coincide in timing with the fall—and subsequent bounce back—in crime. The major observable variable we consider

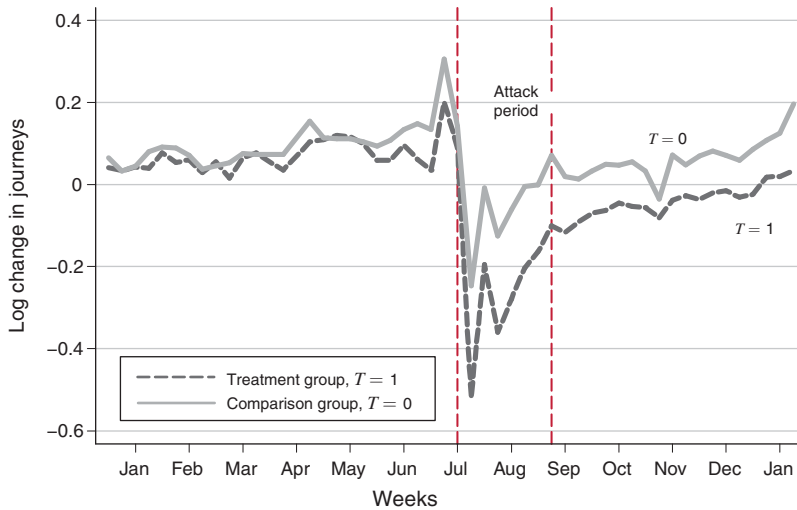


FIGURE 4. YEAR-ON-YEAR CHANGES IN NUMBER OF TUBE JOURNEYS, JANUARY 2004–JANUARY 2006

Notes: Horizontal axis covers the period from January 2005 to January 2006. Note that since this is year-on-year, seasonally differenced data, it reflects an underlying sample extending from January 2004 to January 2006. The vertical axis measures the year-on-year log change in tube journeys. Tube journeys per station are measured as the sum of station entry and exit (i.e., inward and outward journeys) as recorded at station gates. Journeys per station are then aggregated to the borough and treatment/comparison group level for this graph. Data provided by Transport for London (TfL).

here concerns transport decisions and we study this using data on tube journeys obtained from Transport for London. This records journey patterns for the main method of public transport around London and therefore provides a good proxy for shifts in the volume of activity around the city. We aggregate the journey information to borough level and normalize it with respect to the number of tube stations in the borough.

Figure 4 shows how journeys changed year-on-year across the treatment and comparison groups. There is no evidence of a discontinuity in travel patterns corresponding exactly to the timing of the six-week period of increased police presence. In fact, the figure shows a smoother change in tube usage, with the number of journeys trending back up and returning only gradually to pre-attack levels by the end of the year, but with no sharp discontinuity like the police and crime series.

Table 5 formally tests for a difference in journeys across treatment and comparison groups. It shows reduced-form estimates using tube journeys as the dependent variable to test to what extent the fall in tube journeys after the attacks followed the pattern of the police deployment. The estimates indicate that total journeys fell by 22 percent (column 2, controls) over the period of Operation Theseus. Some of this fall, however, may have been due to a diversion of commuters onto other modes of public transport. This is particularly plausible given that two tube lines running through the treatment group were effectively closed down for approximately four weeks after July 7. To examine this, we instead normalize journeys by the number of *open* tube stations with the results reported in panel B of the table. The effect is now smaller, at 13 percent. Importantly, on timing, notice that reduced use of the tube persisted and carried on well after the police numbers had gone back to their original levels.

TABLE 5—CHANGES IN TUBE JOURNEYS, BEFORE AND AFTER JULY 7, 2005

	log(Journeys/number of stations)		log(Journeys/number of open stations)	
	Baseline (1)	Add controls (2)	Baseline (3)	Add controls (4)
$T \times \text{Post-Attack1}$	-0.212*** (0.018)	-0.215*** (0.021)	-0.133*** (0.016)	-0.137*** (0.020)
$T \times \text{Post-Attack2}$	-0.105*** (0.008)	-0.103*** (0.008)	-0.105*** (0.008)	-0.103*** (0.008)
Controls	No	Yes	No	Yes
Observations	104	104	104	104

	log(Journeys/number of stations)		log(Journeys/number of stations)	
	Weekdays (1)	Weekends (2)	Weekdays (3)	Weekends (4)
$T \times \text{Post-Attack1}$	-0.196*** (0.018)	-0.281*** (0.034)	-0.197*** (0.022)	-0.294*** (0.045)
$T \times \text{Post-Attack2}$	-0.097*** (0.010)	-0.106*** (0.032)	-0.093*** (0.010)	-0.112*** (0.030)
Controls	No	No	Yes	Yes
Observations	104	104	104	104

Notes: Borough-level data collapsed by treatment and comparison group, 2 units over 52 weeks. All columns include week fixed effects. Standard errors clustered by treatment group unit in parentheses. All regressions weighted by treatment and comparison group populations. Results adjusted for closed stations (i.e., using the number of open stations as denominator) do not count closed stations along the Piccadilly Line (Arnos Grove to Hyde Park Corner) and Hammersmith and City Line (closed from July 7 to August 2, 2005). Note that stations that intersect with other tube lines are not counted as part of this closure.

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

This final point about the *persistent* effect of the terror attacks on tube-related travel decisions is useful for illustrating the correlated shocks issue. As Table 5 shows, tube travel continued to be significantly lower in the treatment group for the entire period until the end of 2005. For example, columns 2 and 4 show that a persistent 10.3 percent fall in tube travel after the police deployment was completed, which is approximately half of the 21 percent effect seen in the Operation Theseus period. If the change in travel patterns induced by the terrorist attacks was responsible for reducing crime, then we would expect some part of this effect to continue after the deployment.

At this point it is worth reconsidering the week-by-week evidence presented in panels A and B of Figures 3. A unique feature of the Operation Theseus deployment is that it provides us with two discontinuities in police presence, namely, the way that the deployment was discretely switched on and off. The first discontinuity is of course related to the initial attack on July 7. Notably, along with an increased police deployment, this first discontinuity is associated with a similarly timed shift in observable and unobservable factors. In particular, this first discontinuity in police deployment was also accompanied by a similarly acute shift in unobservable factors (that is, widespread changes in behavior and attitudes toward public security risks—“panic” for shorthand). Because these two effects coincide exactly, it is legitimate to raise the argument that the reduction in crime could have been partly driven by the shift in correlated unobservables.

The second discontinuity, however, again provides a useful counterfactual. In this case the police deployment was “switched off” in an environment where unobservable factors were still in effect. Importantly, the Metropolitan Police never made an official public announcement that the police deployment was being significantly reduced. This decision, therefore, limits the scope for unobservable factors to explicitly follow or respond to the police deployment. It is interesting to compare the treatment effect estimates immediately before and after the deployment was switched off. The estimated treatment interaction in week 85 (the last week of the police deployment) was -0.107 (0.043), while the same interaction in the two following weeks is estimated as being -0.040 (0.061) and -0.041 (0.045). This shows that crime in the treatment group increased again at the exact point that the police deployment was withdrawn. Furthermore, this discrete shift in deployment occurred as observable and unobservable factors that could have affected crime still strongly persisted (for example, recall the -10.3 percent gap in tube travel in Table 5 after the deployment was withdrawn).

More generally, this second discontinuity illustrates the point that any correlated, unobservable shocks affecting crime would need to be exactly and exquisitely timed to account for the drop in crime that occurred during Operation Theseus. Our argument, then, is that such timing is implausible given the decentralized nature of the decisions driving changes in unobservables. That is, the unobservable shocks are the result of individual decisions by millions of commuters and members of the public, while Operation Theseus was a centrally determined policy with a clear “on” and “off” date. Indeed, the evidence on the police deployment that we show in this paper indicates that the Metropolitan Police’s response was quite deterministic. That is, deployment levels were raised in the treatment group while carefully keeping levels constant in the comparison group. Furthermore, police deployment levels were effectively restored to their pre-attack levels after Operation Theseus.¹³ In contrast, shifts in travel patterns by inbound commuters did not match the timing and location of the police response.¹⁴

The issue of work travel decisions also uncovers a source of variation that we are able to exploit for evaluating the possible effect of observable, activity-related shocks. Specifically, any basic model of work and nonwork travel decisions predicts interesting variations in terms of timing. For example, we would expect that, faced with the terrorist risks associated with travel on public transport, people would adjust their behavior differently for nonwork travel. That is, the travel decision is less elastic for the travel-to-work decision compared to that for nonwork travel. We would therefore expect that tube journeys would fall by proportionately more on weekends (when most nonwork travel takes place) than on weekdays. This does

¹³ Our discussions with MPS policy officers indicate that big changes in the relative levels of ongoing police deployment in different boroughs occur only rarely. Relative levels of police deployment are determined mainly by centralized formulas (where the main criteria are borough characteristics) with changes determined by a centralized committee.

¹⁴ More support for the hypothesis that changing travel patterns did not match the timing of changes in police presence follows from an analysis of LFS data. These data give information on where people live and work, enabling us to look at whether the number of inbound commuters to inner London changed. There is no evidence that the work travel decisions of people commuting in from outer London and the South East were affected by the attacks in that changes in the proportion of inbound commuters before and after the attacks are statistically insignificant, supporting the idea that modes of transport activity were affected more than travel volumes (see online Appendix Table A4).

TABLE 6—ESTIMATED CRIME TREATMENT EFFECTS WHEN EXCLUDING WEEKENDS

	Reduced form (1)	IV (2)	Reduced form (3)	IV (4)
<i>Panel A. Susceptible crimes</i>			<i>Panel B. Nonsusceptible crimes</i>	
$T \times \text{Post-Attack1}$	-0.138*** (0.046)		0.005 (0.025)	
$T \times \text{Post-Attack2}$	-0.032 (0.031)		-0.037 (0.043)	
ln(police deployment)		-0.400*** (0.150)		0.017 (0.072)
Controls	Yes	Yes	Yes	Yes
Number of boroughs	32	32	32	32
Observations	1,664	1,664	1,664	1,664

Notes: All specifications include week fixed effects. Standard errors clustered by borough in parentheses. Weighted by borough population. These models estimate similar models to Table 4 but use a count of crimes per 1,000 population that excludes all crimes occurring on weekends (i.e., using only Monday–Friday). Treatment groups, T^* *Post-Attack* terms and crime categories defined as in Table 4.

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

seem to have been the case, with tube journeys falling by 28 percent on weekends as compared to 20 percent on weekdays (see the lower panel of Table 5).

Thus, there is an important source of intra-week variation in the shock to observables. If the shock to observables drives the fall in crime, then we would expect this to reflect a more pronounced effect of police on crime on weekends. Following this, we have reestimated the baseline models excluding all observations relating to weekends.¹⁵ This results in very similar coefficient estimates, and only slightly larger standard errors, as shown in Table 6. Importantly, this means that our estimates are unaffected, even when we drop the section of our crime data that is most vulnerable to the problem of correlated observable shocks.

A similar argument can be made in terms of correlated unobservable shocks. As already seen, there is a distinctive pattern to the timing of the fall in crime and its subsequent bounce back. For unobservable shocks to be driving our results, their effect would have to be large and exquisitely timed to perfectly match the police and crime changes. Basic survey evidence on risk attitudes among inner and outer London residents, however, also suggests no significant difference in the types of attitudes that would drive a set of significant, differential unobservable shocks across treatment and comparison groups. Indeed, responses on attitudes to the terror attacks given by inner and outer London residents are closely comparable.¹⁶ The attacks almost certainly had an impact on risk attitudes, but they seem to be very similar in the treatment and comparison areas. From this we conclude that the effect of unobservables is likely to be minimal.

¹⁵ Recall that our crime, police, and tube journeys data are available at daily level for the years 2004–2005. This gives us the flexibility to drop Saturday and Sunday prior to aggregating to a weekly frequency.

¹⁶ See online Appendix Table A5.

F. Possible Crime Displacement

The final empirical issue we consider is that of crime displacement, both spatial and temporal. These two displacement effects have opposing effects on the crime-police relationship we have estimated. First, spatial displacement into comparison areas is likely to impart a downward bias on our estimate, as it will move criminal activity into the nontreated boroughs, increasing crime there and lowering the DiD estimate. Second, temporal displacement could impart an upward bias on our estimate. Criminals operating in the treatment group could delay their actions, thus contributing to a larger fall in crime during the policy-on period, but subsequently there will be a compensating increase in crime in the wake of the policy.

It turns out that we are not able to marshal evidence for either of these.¹⁷ The paper by Draca, Machin, and Witt (2010) looks at spatial displacement due to Operation Theseus in more detail. Here, we note the results of a robustness check where the comparison group is restricted to a set of adjacent and/or central London boroughs. If crime were displaced to these geographically closer boroughs, then we would see different estimates from the baseline estimates considered earlier. In particular, if crime rose in these nearby boroughs as a result of displacement, then we would expect a smaller DiD estimate. Using these more matched boroughs (Adjacent and Central Ten¹⁸) produces very similar results to the earlier estimates. Compared to the baseline estimates discussed earlier (-0.132 , with associated standard error 0.031), for susceptible crimes the estimated effects (standard errors) were -0.129 (0.040) for Adjacent and -0.108 (0.051) for Central Ten. Thus, the estimates are similar, identifying a crime fall of around 11–13 percent for susceptible crimes in central London relative to the (respective) comparison boroughs. In line with the earlier baseline results, there was no impact on nonsusceptible crimes. As such, it does not seem that contemporaneous spatial displacement occurred during Operation Theseus.¹⁹

Finally, the issue of temporal displacement in the treatment group can be looked at by referring back to the week-by-week estimates of treatment effects in Figure 3. As we have already noted, after the bounce-back in treatment borough crime that quickly occurred at the end of Operation Theseus, there is no evidence in the differenced models of any subsequent increase in treatment area crime. This seems to run against the hypothesis of intertemporal substitution where criminal activity rebounds after the police deployment was withdrawn from the treatment boroughs.²⁰

¹⁷ On the face of it, the lack of displacement for the kind of susceptible street-level crimes where we find effects might seem surprising. Unfortunately, we do not have data that would enable a stronger test to be undertaken (e.g., on tourist traffic in treatment and control boroughs, or on the attitudes of criminals).

¹⁸ Adjacent boroughs were: Brent, Hackney, Hammersmith and Fulham, Lambeth, Newham, Southwark, and Wandsworth. Central Ten boroughs were: Westminster, Camden, Islington, Kensington and Chelsea, Tower Hamlets (Treatment Group), and Brent, Hackney, Hammersmith and Fulham, Lambeth, and Southwark.

¹⁹ As a further check for displacement effects, we followed the approach of Jeffrey Grogger (2002) in contrasting crime trends between adjacent and nonadjacent comparison boroughs. Again, however, we could not uncover decisive evidence of between-borough displacement effects.

²⁰ Closer inspection of panel B of Figure 2 does show something of an upturn in crime in the comparison boroughs in the latter part of 2005. Indeed, it is possible that this could reflect a delayed spatial displacement effect. Counter to this, however, in the differenced statistical models, treatment-control borough differences in the post-Operation Theseus period are not statistically significant (see the insignificant coefficients on $T \times \text{Post-Attack2}$ in Tables 2, 3, 4, and 6).

IV. Conclusions

In this paper we provide new, highly robust evidence on the causal impact of police on crime. We find strong evidence that more police lead to reductions in what we refer to as susceptible crimes (i.e., those that are more likely to be prevented by police visibility, including street crimes like robberies and thefts). Our starting point is the basic insight at the center of Di Tella and Schargrodsky's (2004) paper, namely, that terrorist attacks can induce exogenous variations in the allocation of police resources that can be used to estimate the causal impact of police on crime. Using the case of the July 2005 London terror attacks, our paper extends this strategy in two significant ways. First, the scale of the police deployment we consider is much greater than the highly localized responses that have previously been studied. Together with the unique police hours data we use, this allows us to provide new, highly robust IV-based estimates of the crime-police elasticity. Furthermore, there is a novel *ceteris paribus* dimension to the London police deployment. By temporarily extending its resources (primarily through overtime), the police service was able to keep their force levels constant in the comparison group that we consider, while simultaneously increasing the police presence in the treatment group. This provides a clean setting to test the relationship between crime and police.

The research design delivers some striking results. There is clear evidence that the timing and location of falls in susceptible crimes closely coincide with the increase in police deployment. Crime rates quickly returned to pre-attack levels after the six week "policy-on" period. Shocks to observable activity (as measured by tube journey data) cannot account for the timing of the fall, and it is hard to conceive of a pattern of unobservable shocks that could do so. However, as with other papers that adopt a "quasi-experimental" approach, one might have concerns about the study's external validity. Using a very different approach from other papers looking at the causal impact of crime, our preferred IV causal estimate of the crime-police elasticity is approximately -0.38 (for susceptible crimes), which is strikingly similar to existing results in the literature (e.g., Levitt 1997; Corman and Mocan 2000; Di Tella and Schargrodsky 2004). Moreover, because of the scale of the deployment change and the very clear coincident timing in the crime fall, this elasticity is precisely estimated and supportive of the basic economic model of crime, in which more police reduce criminal activity.

REFERENCES

- Becker, Gary, and Yona Rubinstein. 2004. "Fear and the Response to Terrorism: An Economic Analysis." Unpublished.
- Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan. 2004. "How Much Should We Trust Differences-in-Differences Estimates?" *Quarterly Journal of Economics*, 119(1): 249–75.
- Bloom, Nicholas. 2009. "The Impact of Uncertainty Shocks." *Econometrica*, 77(3): 623–85.
- Corman, Hope, and H. Naci Mocan. 2000. "A Time-Series Analysis of Crime, Deterrence, and Drug Abuse in New York City." *American Economic Review*, 90(3): 584–604.
- Di Tella, Rafael, and Ernesto Schargrodsky. 2004. "Do Police Reduce Crime? Estimates Using the Allocation of Police Forces after a Terrorist Attack." *American Economic Review*, 94(1): 115–33.
- Donohue, John J. 1998. "Understanding the Time Path of Crime." *Journal of Criminal Law and Criminology*, 88(4): 1423–52.
- Draca, Mirko, Stephen Machin, and Robert Witt. 2010. "Crime Displacement and Police Interventions: Evidence from London's 'Operation Theseus.'" In *The Economics of Crime: Lessons for and*

- from *Latin America*, ed. Rafael Di Tella, Sebastian Edwards and Ernesto Schargrotsky, 359–78. Chicago: University of Chicago Press.
- Draca, Mirko, Stephen Machin, and Robert Witt.** 2011. “Panic on the Streets of London: Police, Crime, and the July 2005 Terror Attacks: Dataset.” *American Economic Review*. <http://www.aeaweb.org/articles.php?doi=10.1257/aer.101.5.2157>.
- Evans, William N., and Emily G. Owens.** 2007. “COPS and Crime.” *Journal of Public Economics*, 91(1–2): 181–201.
- Gottfredson, Michael R., and Travis Hirschi.** 1990. *A General Theory of Crime*. Stanford: Stanford University Press.
- Greater London Authority.** 2005. *London’s Economic Outlook: Autumn 2005*. London: Greater London Authority.
- Grogger, Jeffrey.** 2002. “The Effects of Civil Gang Injunctions on Reported Violent Crime: Evidence from Los Angeles County.” *Journal of Law and Economics*, 45(1): 69–90.
- Klick, Jonathan, and Alexander Tabarrok.** 2005. “Using Terror Alert Levels to Estimate the Effect of Police on Crime.” *Journal of Law and Economics*, 48(1): 267–79.
- Klockars, Carl.** 1983. *Thinking about Police: Contemporary Readings*. New York: McGraw-Hill.
- Krueger, Alan B.** 2007. *What Makes a Terrorist: Economics and the Roots of Terrorism*. Lionel Robbins Lectures. Princeton, NJ: Princeton University Press.
- Krueger, Alan B., and Jitka Maleckova.** 2003. “Education, Poverty and Terrorism: Is There a Causal Connection?” *Journal of Economic Perspectives*, 17(4): 119–44.
- Lenain, Patrick, Marcos Bonturi, and Vincent Koen.** 2002. “The Economic Consequences of Terrorism.” Organisation for Economic Co-operation and Development Working Paper 334.
- Levitt, Steven D.** 1997. “Using Electoral Cycles in Police Hiring to Estimate the Effect of Police on Crime.” *American Economic Review*, 87(3): 270–90.
- Levitt, Steven D.** 2002. “Using Electoral Cycles in Police Hiring to Estimate the Effects of Police on Crime: Reply.” *American Economic Review*, 92(4): 1244–50.
- London Metropolitan Police Service.** 2006. Computer Aided Resource Management. Data extract provided by LMPS.
- London Metropolitan Police Service.** 2006. Crime Mapping Statistics. Data extract provided by LMPS.
- Machin, Stephen, and Olivier Marie.** Forthcoming. “Crime and Police Resources: The Street Crime Initiative.” *Journal of the European Economic Association*.
- McCrary, Justin.** 2002. “Using Electoral Cycles in Police Hiring to Estimate the Effect of Police on Crime: Comment.” *American Economic Review*, 92(4): 1236–43.
- Office for National Statistics. Social and Vital Statistics Division and Northern Ireland Statistics and Research Agency. Central Survey Unit.** 2006. Labour Force Survey 2004–2006 Colchester, Essex: UK Data Archive.
- Office for National Statistics. Population Estimates Borough.** 2004–2006. <http://data.london.gov.uk/datastore/package/office-national-statistics-ons-population-estimates-borough> (accessed June 29, 2011).
- Sherman, Lawrence W., and David Weisburd.** 1995. “General Deterrent Effects of Police Patrols in Crime ‘Hot Spots’: A Randomized Controlled Trial.” *Justice Quarterly*, 12(4): 625–48.
- Transport for London.** 2006. Tube Journeys Statistics. Data extract provided by Transport for London.