Does Terrorism Increase Crime? A Cautionary Tale*

John J. Donohue III†
Yale Law School

Daniel E. Ho‡
Institute for Quantitative Social Science, Harvard University
Yale Law School

First draft, August 8, 2005
This draft, August 8, 2005

Abstract

How does the war on terrorism affect crime? We investigate a natural experiment first studied by Di Tella and Schargrodsky (2004), which found that increased police presence in certain blocks due to a terrorist attack decreases crime. Our examination suggests that the story may be more complex. First, our findings confirm that increased police presence appears to reduce crime in those blocks that received 24-hour surveillance due to a terrorist attack (protected blocks). However, increased police surveillance may also have *increased* crime in unprotected blocks by (a) shifting police forces away from or (b) displacing ordinary criminal activity to unprotected blocks. Crime levels increase substantially in unprotected districts that are located several blocks away from protected ones. This effect appears to be larger than that of crime deterred, suggesting that shifting police resources to fight terrorism may in fact increase absolute levels of ordinary crime.

1 Terrorism, Police, and Crime

Understanding the causal effect of police presence on crime remains a long-standing puzzle in the economics of criminal behavior (Becker, 1968; Cameron, 1988). Credible answers to the question of whether police presence deters crime are difficult to obtain because of simultaneity of police levels and crime (Levitt, 1997; McCrary, 2002; Levitt, 2002; Marvell and Carlisle, 1996; Corman and Mocan, 2000). To break this simultaneity, Di Tella and Schargrodsky (2004) ("DS") examined

*We are grateful to Rafael Di Tella and Ernesto Schargrodsky for making their data available. Research support was provided by the Institute for Quantitative Social Science, Harvard University.

†Leighton Homer Surbeck Professor of Law; Yale Law School, P.O. Box 208215, New Haven, CT 06520, Email: j.donohue@Yale.edu

‡Ph.D., Department of Government, Harvard University; J.D. Candidate, Yale Law School; Phone 617–642–5904, Fax: 617-496-2254, Email: daniel.ho@yale.edu, URL: people.iq.harvard.edu/~dho

a natural experiment to isolate the causal effect of police presence on crime. In particular, DS capitalized on a terrorist attack on the main Jewish center in Buenos Aires in 1994. In response to this horrendous attack, which killed 85 and wounded 300, the federal government assigned 24-hour police protection to every Jewish and Muslim institution in the country. Employing a difference-in-differences ("DID") approach, DS compared auto thefts in two locales – (1) neighborhood blocks before and after assignment of 24-hour police protection ("protected blocks"), and (2) blocks that were not assigned this special protection ("unprotected blocks"). This approach bears substantial promise since the temporal and geographic distribution of additional police protection can be assumed to be plausibly exogenous (see also Klick and Tabarrok, 2005). DS found "a large deterrent effect on observable police crime...[and that] [t]he effect is local, with no appreciable impact outside the narrow area in which the police deployed" (p. 115).

Here we build on this important work to examine the effects of terrorism and police reallocation on both protected and unprotected blocks. In particular, we investigate whether there is any evidence for a "displacement effect" – that is, increased crime in the unprotected blocks that would bias upward the DS estimates. DS acknowledged that "it is still possible that car thefts were displaced in a way that we are unable to measure, in which case the effect of policing may be smaller than our estimates suggest" (p. 117). Using the same time-series approach that DS used for protected blocks, we find that unprotected blocks that served as a control group in the DS analysis appear to experience substantial increases in crime after the terrorist attack. This suggests that the independence assumption of the DID estimator may be violated, since police reallocation to protected blocks may affect unprotected blocks. Substantively, police forces may have been shifted away from unprotected areas as a result of the terrorist attack, and/or added police protection in protected blocks may have displaced crime to unprotected blocks. Providing 24-hour protection to some areas to fight terrorism appears to require some sacrifice in protecting other areas from ordinary crime. Our evidence suggests that this reallocation of the police force increased aggregate auto thefts, and that this impact was anything but local.

This paper proceeds as follows. Section 2 begins our investigation of the displacement hypothesis with three approaches: exploiting time-series and cross-sectional variation, cross-validation with a pretreatment sample, and examining where crime may have gotten displaced as a result of the police reallocation. We also provide a nonparametric randomization inference approach (Rosenbaum, 2002; Ho and Imai, 2004; Imbens and Rosenbaum, 2005) that addresses well-known variance estimation problems in DID and before-after designs (see also Bertand *et al.*, 2004). These four pieces of evidence are largely consistent with the displacement hypothesis. Section 3 concludes.

2 Evidence for the Displacement Hypothesis

2.1 The DID Approach

To identify the causal effect of police presence on car theft, DS employed a DID estimator:

$$y_{it} = \alpha_0 T_{it}^0 + \alpha_1 T_{it}^1 + \alpha_2 T_{it}^2 + M_t + F_i + \epsilon_{it}, \quad (1)$$

where $T_{it}^j = 1$ denotes increased police presence which equals 1 if block $i = 1, 2, ..., 876$ was j block(s) away from a Jewish or Muslim institution and if the month $t = 4, ..., 12$ was after July (i.e., $t > 7$), M_t represents time fixed-effects and F_i represents block fixed-effects. Two key assumptions to this estimator are that (a) the control group of unprotected blocks is unaffected by the added police protection to protected blocks (independence), and (b) that protected and unprotected blocks follow the same time trend. Figure 1 plots average monthly car thefts per block in the three Buenos Aires neighborhoods studied by DS from April to December 1994. The top panel presents blocks that would receive added police protection in August, and the middle panel presents blocks that serve as the control group. While this figure shows the protected blocks appear to have experienced declines in car thefts, the middle panel also suggests that crime increases in unprotected blocks after the terrorist attack. The bottom panel depicts averages over time. Mean car theft levels are indistinguishable between protected and unprotected blocks before the attack, but after the attack the mean level for protected blocks decreases while it increases for unprotected blocks.

Under the DID assumption, the overall trend for the control group would present the "natural" occurrence of car thefts among protected institutions but for the terrorist attack. This could be implausible if there is block-specific seasonality in crime (which implies there is no natural benchmark for crime trends in the protected blocks) or if independence is violated due to police or crime displacement. This is far from a minor technical concern. Indeed, DS expressly acknowledged the possibility of displacement (p. 117) and anecdotal evidence suggests that substantial police reallocation occurred. As DS noted, for example, roughly one third of police officers in the neighborhood of Once had to be reassigned to protection duties (p. 117). That said, DS posited that police displacement may not be a large issue because reassignment may have been largely from administrative tasks and from outside of the districts examined. Of course this is a question that should be subjected to empirical investigation, not assumption. Moreover, even if police reallocation had no

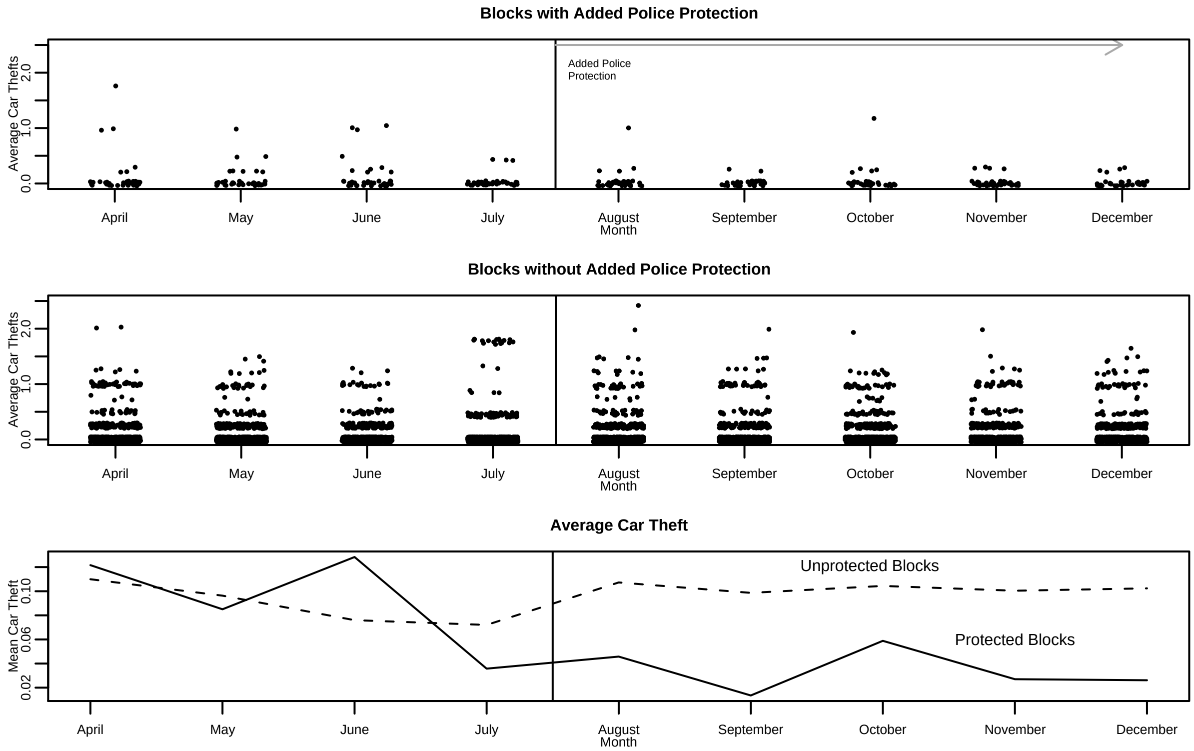

Figure 1: Average Monthly Car Thefts Per Block from April to December of 1994 in Three Buenos Aires Neighborhoods. In response to a terrorist attack in on July 18, the federal government provided 24-hour police protection starting in August for Jewish and Muslim institutions, the blocks of which are depicted in the top panel. The middle panel depicts blocks that received no additional police protection. The bottom panel depicts mean trends for protected and unprotected blocks. This figure shows that while there was a downward trend in car thefts prior to the terrorist attack, car thefts appear to have *increased* in unprotected blocks in August, indicating that the DID assumption may be violated by crime or police displacement. Observations are randomly jittered for visibility. Car thefts occurring between July 18 and July 31 are excluded as in DS – thefts are therefore normalized to 30-day months.

effect on police levels in control blocks, this reallocation could still have displaced criminal activity to adjacent blocks. In sum, DS's natural experiment would be perfect if: (1) police staffing in the unprotected areas had been unchanged after the terrorist attack as the staffing in the protected area rose with police of average quality; and (2) crime in the protected and unprotected blocks was not linked, so that the higher police protection in one block did not simply shift crime to the unprotected blocks. The danger that displacement of criminals would contaminate the control group is substantial. While DS conducted an exemplary slew of sensitivity and robustness analyses, we focus here on the implications of this displacement hypothesis, which was not assessed in the original article.

2.2 More Evidence

In addition to their DID analysis, DS also presented time-series evidence for blocks close to the protected institutions that showed a significant drop in car thefts on the protected block after the police reassignment (DS, Table 3E). In Table 1, we employ a similar time-series approach, but unlike DS, we apply it only to unprotected blocks. Our time-series estimation confirms the trends suggested by Figure 1: Column (A) suggests that unprotected blocks experience an increase in roughly 0.014 car thefts per month, representing a 13 percent increase from before the terrorist attack.[1] Columns (B) and (C) test whether this effect may be driven by spillover effects from protected blocks, estimating the effect on blocks that are one or two blocks away from Jewish or Muslim institutions. There is a hint of evidence (albeit not statistically significant) that the protective halo from the increased police presence extends for one block. Beyond that, the control for nearby blocks leads to slightly higher estimates of the spillover crime increase in the remaining unprotected blocks.

If police forces were optimally allocated to fight crime prior to the terrorist attack, we might expect that reallocation of police forces to protect Jewish and Muslim institutions would have a net negative effect on ordinary crime. Since only 37 of 876 blocks (roughly 4 percent) contain a Jewish or Muslim institution, even a small increase in crime in these districts may overwhelm the estimated auto theft decrease of roughly 0.081 in protected blocks. Column (D) confirms this, by testing whether aggregate auto theft levels were greater before or after the attack. Across all districts there was an increase of 0.011 car thefts per month per block after the terrorist attack.

Relying on such simple pre-post comparisons can be problematic. Most importantly, serial correlation can lead to understated standard errors, which may in turn lead us to falsely attribute auto theft changes to police reallocation. Such serial correlation, which is well-known to plague DID estimators (Bertand *et al.*, 2004), may be particularly acute given the seasonality of auto theft that is typical in the United States (Corman and Mocan, 2000). If auto theft in Argentina similarly decreases in the winter season, simple time-series estimates may overestimate the deterrent effect of police while underestimating the displacement effect (recall that the peak of winter in Argentina occurs roughly in July-August). DS conducted a host of specification tests by primarily reestimating standard errors under a series of clustering assumptions. To assess sensitivity here, we employ randomization inference (Fisher, 1935; Rosenbaum, 2002; Ho and Imai, 2004), which

[1]This is calculated from the baseline average number of car thefts for all blocks more than two blocks away from a protected institution of 0.108 (i.e., 0.014/0.108 = 0.130).

	Unprotected Blocks Only			All Blocks
	(A)	(B)	(C)	(D)
Post Terrorist Attack	0.0140**	0.0174***	0.0176**	0.0110**
	(0.0055)	(0.0060)	(0.0074)	(0.0054)
One-Block Police & Post Terrorist Attack		−0.0174	−0.0176	
		(0.0151)	(0.0157)	
Two-Block Police & Post Terrorist Attack			0.0006	
			(0.0128)	
Block fixed effect	Yes	Yes	Yes	Yes
Month fixed effect	No	No	No	No
Number of observations	7551	7551	7551	7884
R^2	0.1895	0.1897	0.1897	0.1896

Table 1: Least-Squares Dummy Variables Regression of Average Car Thefts Per Month Per Block Employing Same Time-Series Approach of DS (p. 123, Column (E)). The positive findings on Post Terrorist Attack indicate that car thefts increased in unprotected blocks as police forces were reallocated. One-Block/Two-Block Police & Post Terrorist Attack indicates whether a unit was one or two blocks away from a protected institution and observed after the terrorist attack. When no month fixed effects are used, the dependent variable is normalized to a 30-day month as in DS (p. 124). Columns (A)-(C) exclude 333 protected blocks. Robust standard errors in parentheses. ** (***) indicates significance at the 5-percent (1-percent) level.

capitalizes on the assumption that the timing of the terrorist attack is exogenous. Randomization inference tests the sharp null hypothesis of no in-sample effects by treating time of the terrorist attack as the only random variable. Using all permutations of the treatment yields an exact randomization distribution, which asymptotic distributions are designed to approximate under certain assumptions (Imbens and Rosenbaum, 2005). This has a primary benefit of not relying on distributional assumptions about the error term, thereby addressing the same types of concerns raised by Bertand *et al.* (2004) and Moulton (1990).

As a first cut, for the full sample we activate time dummies at each of the 8 months of the dataset. This reveals that the coefficient with the highest effect is when July is treated as the breaking point. Given the low number of months, however, we have little power to reject the null hypothesis of no impact of terrorism (the lowest p-value, which we obtain here, is $1/8 = 0.125$). We hence conduct the same analysis activating weekly time dummies to compare crime levels before and after any given week. This yields 36 possible values and thereby higher power (the lowest p-value would be $1/36 = 0.03$). Figure 2 plots all 36 test statistics, which simply represent difference in mean car theft levels before and after. The solid dots indicate test statistics for unprotected blocks only. Under the null hypothesis of no effects, the observed test statistic of 0.0024 car thefts per week is the third highest in the sample (p-value=0.08). For protected blocks, on the other hand,

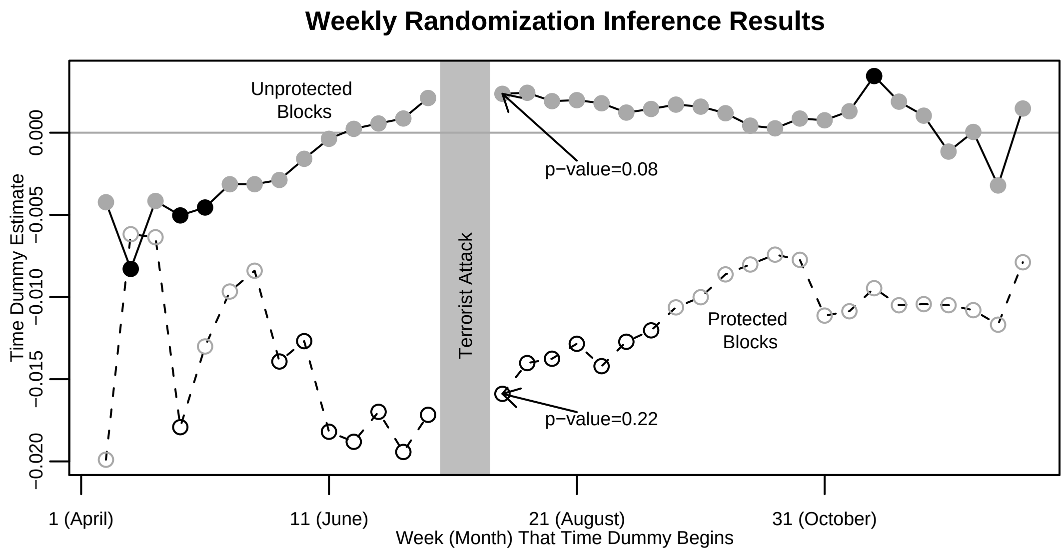

Figure 2: Plot of Randomization Distribution of Moving Weekly Time Dummies. The x-axis depicts at which week a time dummy is activated for a total of 36 weeks, excluding car thefts occurring between July 18 and July 31. The y-axis depicts the least squares estimate of the time dummy coefficient. Solid (hollow) dots indicate estimates for unprotected (protected) blocks. Dots are black to indicate statistical significance if the absolute t-statistic exceeds 2, and grey otherwise. This figure shows that using time-series randomization inference we find that the police reallocation due to the terrorist attack increased crime in unprotected blocks (p-value=0.08), but had a negligible impact on protected blocks (p-value=0.22).

we cannot reject the null hypothesis of no effects (p-value=0.22). Using the full sample, the p-value for the observed test statistic of an increase of 0.0016 car thefts per week is 0.08, suggesting a positive overall effect of the terrorist attack on car theft. This exercise reveals an important virtue of randomization inference compared to conventional estimators, akin to that discussed by Bertand *et al.* (2004): for protected blocks, we would estimate a significant effect for over 41% of the permutations at a 0.05 level, even if we have randomly selected the activation of time dummies. These apparent false positives are indicated by the black (versus grey) dots in Figure 2, 15 of which are significant for the protected blocks depicted on the lower part of the graph.

Randomization inference can also incorporate the additional exogeneity assumption regarding the distribution of blocks to which additional police protection was assigned, thereby gaining larger power depending on the test statistic. Specifically, assuming that a terrorist attack could have targeted some institution such that police protection would have been provided to any other combination of 37 other blocks, we have a total of $8 \times \binom{876}{37}$ possible permutations. Since this number exceeds 2.0×10^{66}, we approximate this distribution with Monte Carlo sampling. (Of course, it is questionable whether terrorist attacks, which usually focus on highly salient public

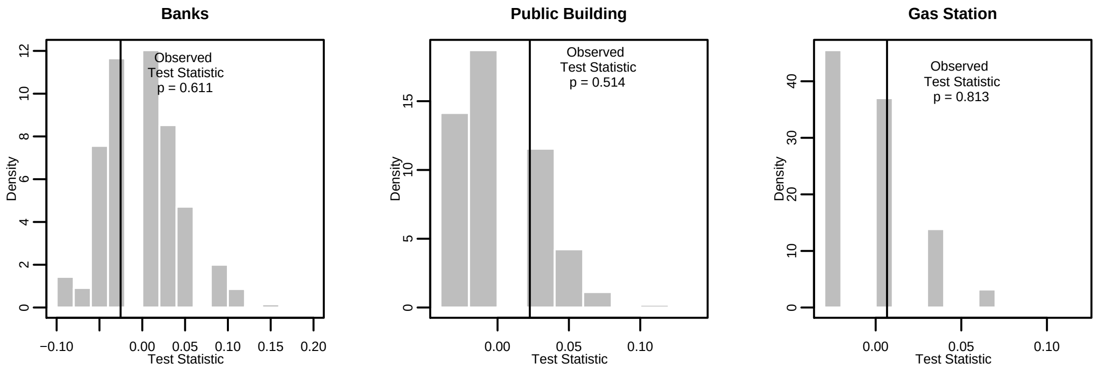

Figure 3: Randomization Inference for Exogenous Indicator Variables of Presence of Bank, Public Building, and Gas Station in Block. This figure shows that assuming random timing of the terrorist attack and random targeting of a subset of 37 of 876 blocks yields no detectable difference between protected and unprotected blocks. Randomization distribution is approximated with 10,000 Monte Carlo simulations.

structures, could truly select from a population of targets that yields such a high number of blocks. If the assumption is not met, a researcher may need to redefine the potential target population, which may decrease the power of the test.) Testing the effects of a time dummy interacted with a dummy on protected and unprotected districts yields p-values of 0.005 and 0.000 for a decrease of 0.07 and an increase of 0.02, respectively, with 10,000 simulations. The former is consistent with the overall deterrence findings of DS, while the latter suggests an additional dimension consistent with the displacement hypothesis.[2]

To compare the effects, Figure 3 presents the randomization distribution for exogenous covariates of the presence of a bank, public building, or gas station on a block. For each of these the observed statistic is well within the mid-range of the simulated randomization distribution, meaning that there are no differences along protected and unprotected blocks before and after the actual terrorist attack.

Our above investigation provides robust evidence for the displacement hypothesis, suggesting that the DID assumptions are violated. Nonetheless, our findings suggest that the DS story appears to be correct, in that "a posted and visible police guard" dampens crime locally, but incorrect in asserting that this effect is "large, negative, [and] local [with] little or no effect outside a narrow area" (p. 131). Due to the terrorist attack, police forces and crime appear to have been displaced, thereby leading car thefts to increase in unprotected districts.

[2]Randomization inference can of course also incorporate the original DID test statistic, yielding a p-value of 0.05. To obtain fully nonparametric confidence intervals, researchers can additionally invert the test as shown by Imbens and Rosenbaum (2005) and Ho and Imai (2004).

	(A)	(B)	(C)
Same-Block Police	0.0080	0.0080	0.0080
	(0.0213)	(0.0214)	(0.0215)
One-Block Police		−0.0114	−0.0114
		(0.0099)	(0.0102)
Two-Block Police			−0.0322
			(0.0091)
Block fixed effect	No	No	No
Month fixed effect	Yes	Yes	Yes
Number of observations	3504	3504	3504
R^2	0.0147	0.0147	0.0147

Table 2: Cross-Section Least Squares Dummy Variable Regressions of Average Car Thefts Per Month Per Block on All Blocks Before Terrorist Attack. These null results are consistent with the idea that systematic differences between protected and unprotected blocks do not drive any findings. Robust standard errors in parentheses. ** (***) indicates significance at the 5-percent (1-percent) level.

2.3 Cross-Validation with Pretreatment Sample

To further validate our hypothesis, we conduct a series of robustness tests with the pretreatment sample only. First, we investigate whether we can detect cross-sectional differences in protected and unprotected blocks. Figure 2 presents results estimating effects on same-block police blocks and blocks that are one or two blocks away from protected blocks. If we can detect systematic differences in these blocks prior to the terrorist attack, this would call into question the DID approaches use of these blocks as a control group. Each of the estimates is indistinguishable from 0, however, suggesting that prior findings are not driven by preexisting differences between protected and unprotected blocks. This is consistent with the other randomization checks conducted in DS (p. 131).

Second, Table 3 examines the time dynamics of the pretreatment sample, by activating police dummies on any of the three months prior to the actual terrorist attack. This is similar in spirit to the randomization approach presented above, and DS conducted an analogous robustness check using the DID specification. These results suggest caution with simple time-series approaches. Both protected and unprotected blocks exhibited sharp declines in car theft rates even prior to the terrorist attack. Columns (A)-(C) present sharply negative effects for all pretreatment blocks, which remain at the same level when excluding protected blocks in columns (D)-(F). We cannot reject the null hypothesis that these downward trends are the same for both protected and unprotected pretreatment blocks. The bottom panel shows that this is also the case when interacting the treatment with protected block dummies, although some of the direct downward trend is soaked up by the interactions when April 30 serves as the cutoff. This Table provides one justification for why randomization inference approaches that make no assumptions about the time dynamics may be preferable to this parametric approach.

2.4 Testing Displacement Effects

If the displacement hypothesis is correct, we should see a clear geographic pattern in crime rates before and after the terrorist attack. Specifically, we might predict a weakly monotonic increase in crime levels in blocks adjacent to protected blocks. To verify this, Figure 4 depicts a contour plot of density of crime levels on the y-axis and distance to Jewish or Muslim institution on the x-axis (0 indicating a protected block). The left panel depicts the density of car theft levels prior to the terrorist attack: the contour line of 0.005, for example, indicates that the highest month-block rate of 0.4 car thefts occurred in blocks that were 1-4 blocks away from a Jewish/Muslim institution. The right panel depicts the same contour lines for the months after the terrorist attack. This depicts a substantial shift of crime away from protected blocks. Contour lines all intersected the left hand axis at higher levels in the left panel. Moreover, each of the contour lines in unprotected blocks shifts outward, most notably for blocks that are 2-3 blocks away from protected blocks. Every level of month-block car theft increases substantially. For example, prior to the terrorist attack roughly 0.026 of all month-blocks recorded a 0.25 average car thefts when 2 blocks away from protected blocks. That number climbs from 0.026 to 0.044 after the terrorist attack.

To test this more formally, we estimate equations with a dummy variable for post-terrorist attack and interacted this variable with indicators for how many blocks away from a protected institution the block was. Figure 5 plots the point estimates in the white line and 95% confidence intervals, with the estimated impact on car thefts on the y-axis and distance on the x-axis. This figure depicts the same story of the contour plot: the terrorism attack has a negative impact on protected blocks, but shifts car thefts upward in virtually all unprotected blocks, most notably in blocks that are just 2-4 blocks away from protected blocks. This suggests that there may be a substantial redistributive component to fighting terrorism and police reallocation. While reallocation has substantial benefits in decreasing ordinary crime for protected blocks, it increases crime in proximate blocks.

	Raw Treatment Effect					
	All Pretreatment Blocks			Unprotected Pretreatment Blocks		
	Police dummies activated on					
	April 30	May 31	June 30	April 30	May 31	June 30
	(A)	(B)	(C)	(D)	(E)	(F)
Treatment	−0.0289***	−0.0288***	−0.0243**	−0.0285***	−0.0291***	−0.022**
	(0.0097)	(0.0080)	(0.0098)	(0.0099)	(0.0081)	(0.0101)
Block fixed effect	Yes	Yes	Yes	Yes	Yes	Yes
Month fixed effect	No	No	No	No	No	No
Number of observations	3504	3504	3504	3356	3356	3356
R^2	0.3105	0.3113	0.3097	0.3089	0.3099	0.3079
	Treatment with Block Interactions					
	All Pretreatment Blocks					
	Police dummies activated on					
	April 30	May 31	June 30	April 30	May 31	June 30
	(G)	(H)	(I)	(J)	(K)	(L)
Treatment	−0.0285***	−0.0291***	−0.0220**	−0.0153	−0.0283***	−0.0250**
	(0.0099)	(0.0081)	(0.0100)	(0.0128)	(0.0105)	(0.0127)
Same-Block Police	−0.0101	0.0079	−0.0539	−0.0233	0.0071	−0.0509
	(0.0536)	(0.0406)	(0.0406)	(0.0542)	(0.0411)	(0.0413)
One-Block Police				−0.0202	0.0211	0.0379
				(0.0274)	(0.0235)	(0.0319)
Two-Blocks Police				−0.0346	−0.0180	−0.0159
				(0.0234)	(0.0189)	(0.0217)
Block fixed effect	Yes	Yes	Yes	Yes	Yes	Yes
Month fixed effect	No	No	No	No	No	No
Number of observations	3504	3504	3504	3504	3504	3504
R^2	0.3105	0.3113	0.3101	0.3112	0.312	0.3111

Table 3: Least-Squares Dummy Variables Regression of Average Car Thefts Per Month Per Block Examining Time Effects in Pretreatment Sample. The top panel shows that there is a substantial downward trend in car thefts prior to the terrorist attack. We cannot reject the hypothesis that the downward trend is the same for both protected and unprotected blocks. The bottom panel shows that some of the downward trend is soaked up by protected blocks and blocks that are several blocks away from protected institutions. This table demonstrates the crucial role in modeling baseline time trends in crime rates for identifying causal effects by DID. When no month fixed effects are used, the dependent variable is normalized to a 30-day month as in DS (p. 124). ** (***) indicates significance at the 5-percent (1-percent) level.

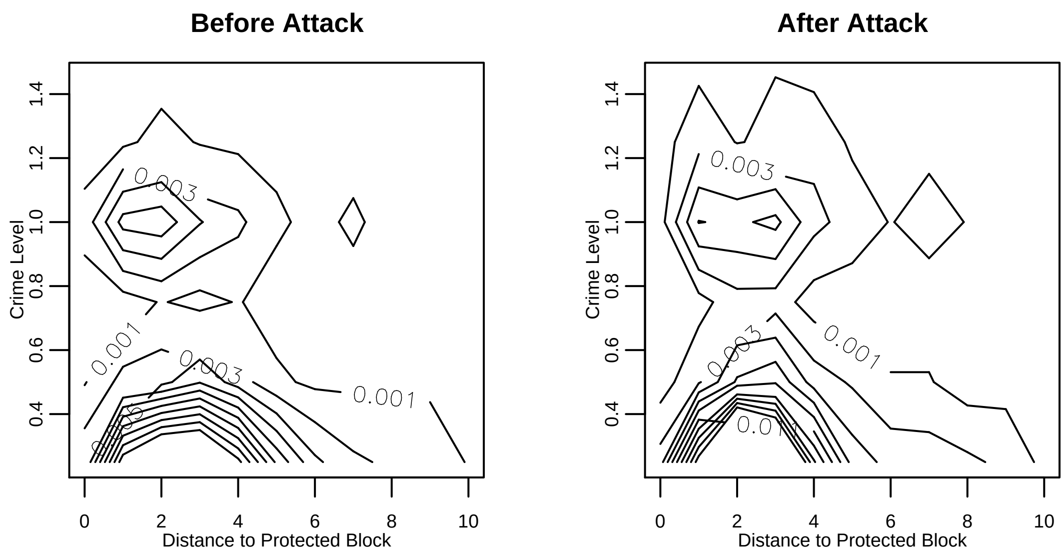

Figure 4: Contour Plot of Car Theft Density by Proximity to Protected Block. The density is estimated by the proportion of car thefts occurring at a particular distance before or after the terrorist attack. This figure shows that crime levels were substantially higher in protected blocks (distance equals 0) but that these crimes appear to be displaced to unprotected blocks that are several blocks away. Unprotected blocks just two or three blocks away from a Jewish or Muslim institution, for example, experience sharp increases in crime levels denoted by the upward movement of the contour lines to above 0.4.

3 Conclusion

To summarize what we have learned in our investigation, Table 4 characterizes four potential states depending on the nature of displacement effects. The columns distinguish whether unprotected blocks are hermetically sealed; that is, whether added police protection in protected blocks displaces crime towards unprotected blocks. The rows distinguish whether the additional protection constituted a pure increase of police in protected blocks or whether police forces were shifted from unprotected to protected blocks. DS assumed the conditions of Cell (1) in the upper left: namely, (a) that no crime is displaced to unprotected blocks, and (b) that the increased police presence constituted a pure increase in resources. Under these strong conditions, the DID estimates are unbiased. Cell (2) in the upper right indicates that crime is displaced but that the pure increase in police presence in protected blocks is not offset by any reduced policing of unprotected blocks. One would expect some crime shifting, so the DID estimate is biased upward. Cell (3) in the lower left indicates that blocks are hermetically sealed, but that there is some drop in police presence or effectiveness in unprotected areas. We would expect that crime would fall in protected blocks and increase in unprotected blocks. The DID estimate again would overstate the benefit of police.

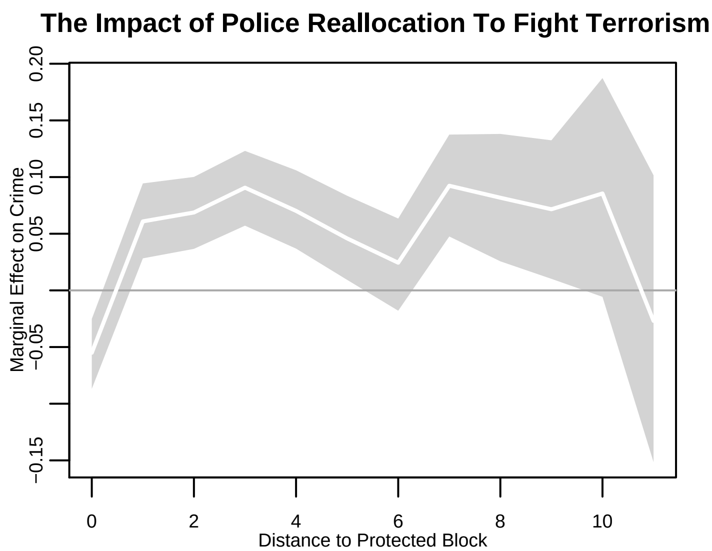

Figure 5: The Marginal Impact of Police Reallocation on Blocks Estimated from a Regression of Average Car Theft Per Month Per Block on a Dummy Variable for Post-Terrorism and this Dummy Interacted with Distance Dummies. White line indicates point estimate and grey bands indicate 90% confidence intervals from simulated posterior distribution. This figure shows that while police reallocation appears to decrease crime levels in protected blocks, reallocation appears to increase in unprotected blocks, particularly in blocks that are 2-4 blocks away from Jewish / Muslim institutions. Since no month fixed effects are used, the dependent variable is normalized to a 30-day month as in DS (p. 124).

Lastly, Cell (4) indicates that crime is displaced to unprotected blocks and that there is a drop in police in unprotected blocks. As a result, we again expect that crime would fall in protected blocks, while increasing in unprotected blocks. The DID estimator will again overstate the benefit of police.

Our analysis confirms that crime fell where more police were stationed for 24-hours a day, but that crime *increased* overall. This is not consistent with efficient criminal conduct under the pure test in Cell (1) because adding police in some blocks while not reducing them elsewhere should not cause an increase in crime (even though it might cause some crime spillovers from the protected blocks). On the other hand, if the number of police did not change, but is only reallocated, it should not be surprising that the areas that lost police protection would experience crime increases and those that received greater protection would experience less crime.

Our analysis shows that the implicit DS assumption that Cell (1) characterizes the Buenos Aires data is incorrect, and that therefore their paper does not succeed in its mission of identifying "the causal effect of police presence on car thefts" (p. 121). With the existing data, however, we cannot go much farther in providing a definitive conclusion about which of the other three cells

	No Crime Displacement	Crime Displacement
Pure Police Increase	(1) DID unbiased Expectation: Auto theft decreases in protected blocks only	(2) DID biased upward Expectation: Auto theft decreases in protected blocks and increases in unprotected blocks
Police Shifted	(3) DID biased upward Expectation: Auto theft decreases in protected blocks and increases in unprotected blocks	(4) DID biased upward Expectation: Auto theft decreases in protected blocks and increases in unprotected blocks

Table 4: The Appropriateness of DID Estimation and Expected Effects on Crime under Four States.

properly characterizes the evidence in Buenos Aires. Several pieces of information would help to shed more light on the precise displacement channel. First, was the increased police presence a pure increase in police or a shifting of police resources? The nature of reassignment would help to assess whether police presence or effectiveness decreased in control blocks. Second, if police were in fact shifted, were they drawn from proximate blocks? If police were shifted from across all districts, our finding of the geographic proximity of crime displacement would suggest that blocks are not hermetically sealed. Third, where are criminals from and what draws them to a particular location? Such information, in conjunction with the geographic displacement patterns, may further enable us to distinguish which of the remaining three states of Table 4 properly characterize the effect of police on crime.

In sum, our findings suggest that fighting terrorism involves complex tradeoffs in police force allocation. Added police protection does appear to reduce crime in protected blocks, but it also appears to increase crime in proximate blocks. Moreover, ordinary levels of crime appear to increase at an aggregate level, suggesting that contrary to DS, the effects are anything but local. Methodologically, our investigation lends further strengths to recent critiques of DID approaches (Bertand *et al.*, 2004). Nonetheless, we offer a number of ways in which scholars can at least partially subject these assumptions to empirical validation.

References

Becker, G. S. (1968). Crime and punishment: An economic approach. *Journal of Political Economy* **76**, 2, 169–217.

Bertand, M., Duflo, E., and Mullainathan, S. (2004). How much should we trust differences-in-differences estimates? *Quarterly Journal of Economics* **119**, 249–75.

Cameron, S. (1988). The economics of crime deterrence: A survey of theory and evidence. *Kyklos* **41**, 2, 301–23.

Corman, H. and Mocan, H. N. (2000). A time-series analysis of crime, deterrence, and drug abuse in new york city. *American Economic Review* **90**, 3, 584–604.

Di Tella, R. and Schargrodsky, E. (2004). Do police reduce crime? estimates using the allocation of police forces after a terrorist attack. *American Economic Review* **94**, 115–133.

Fisher, R. A. (1935). *The Design of Experiments*. Oliver and Boyd, London.

Ho, D. E. and Imai, K. (2004). Randomization inference with natural experiments: An analysis of ballot effects in the 2003 california recall election. *Submitted to Journal of the American Statistical Association* available at `http://www.princeton.edu/~kimai/research/fisher.html`.

Imbens, G. W. and Rosenbaum, P. R. (2005). Robust, accurate confidence intervals with a weak instrument: Quarter of birth and education. *Journal of the Royal Statistical Society, Series A* **168**, 1, 109–126.

Klick, J. and Tabarrok, A. (2005). Using terror alert levels to estimate the effect of police on crime. *Journal of Law and Economics* **48**, 267–279.

Levitt, S. D. (1997). Using electoral cycles in police hiring to estimate the effect of police on crime. *American Economic Review* **87**, 270–90.

Levitt, S. D. (2002). Using electoral cycles in police hiring to estimate the effects of police on crime: Reply. *American Economic Review* **92**, 4, 1244–1250.

Marvell, T. and Carlisle, M. (1996). Specification problems, police levels, and crime rates. *Criminology* **34**, 609–46.

McCrary, J. (2002). Using electoral cycles in police hiring to estimate the effect of police on crime: Comment. *American Economic Review* **92**, 4, 1236–1243.

Moulton, B. R. (1990). An illustration of a pitfall in estimating the effects of aggregate variables on micro unit. *The Review of Economics and Statistics* **72**, 2, 334–38.

Rosenbaum, P. R. (2002). *Observational Studies*. Springer-Verlag, New York.