Police Patrols and Crime

BY JORDI BLANES I VIDAL, LONDON SCHOOL OF ECONOMICS, AND
GIOVANNI MASTROBUONI, COLLEGIO CARLO ALBERTO

The causal relationship between police street deployment and crime represents one of the most common tests of the deterrence mechanism underlying the standard economic model of crime. It is also an important policy parameter in its own right. In England, for instance, foot patrol was at the core of Robert Peel's pioneering vision of a modern police force, and this perspective retains enormous contemporary importance. In the United States, 68 percent of police officers are assigned to patrol operations. Understanding their effect on crime prevention is critical to an accurate evaluation of whether allocating this large resource of patrols is worthwhile, or whether more resources should be allocated to other police operations.

Two previous studies, one by Rafael Di Tella and Ernesto Schargrodsky, and one by Mirko Draca, Stephen Machin, and Robert Witt, estimate the effects on crime of large and sustained increases in static police deployment following deadly terrorist attacks. These two terrorism-based studies identify large and similar elasticities of around 35 percent. Together with the seemingly modest effects of lengthier prison sentences, these findings have led to the demand for visible police deployment to be a central policy lever in terms of crime deterrence.

We regard the terrorism-based studies as persuasive in the claims that deterrence is at work in their specific circumstances.

However, from a policy perspective it is unclear whether these results are generalizable to more common policing strategies. The large, sustained, and concentrated deployments studied in these papers create ideal conditions for police presence to be highly salient to potential offenders. This salience is likely compounded by the fact that the aftermath of a deadly terrorist attack is an occasion when citizens (including potential offenders) are unusually aware of police levels. Both features imply that the elasticities following terrorist attacks may be larger than in other periods.

In general, round-the-clock protection of a highly sensitized citizenry is economically and politically unfeasible. Instead, police protection typically consists of officers moving around large areas while spending little time in each location. Policymakers must therefore evaluate whether the rather thin cover afforded by these patrolling officers is a worthwhile use of resources. The use of terrorist attacks, while useful in terms of exogenous variation, may hinder the extrapolation of the resulting elasticities to these more common policing strategies.

Estimating the effect of police patrols under more normal circumstances is empirically challenging. In addition to the obvious difficulty of isolating exogenous variation in police presence, an appropriate research strategy must leverage sufficient statistical power to identify the effects of potentially small levels of police presence. This requires an exceptionally rich dataset with both a large number of observations and the measurement of police presence with a high level of granularity.

We estimate the effect of police patrols on crime by exploiting a unique dataset tracking police officers in real time, and a policy experiment that created the type of

variation (small, short-lived, and unrelated to extraordinary events) in street deployment that closely resembles typical police patrolling. The elasticity that we find differs from those in the aforementioned terrorism-based papers.

In the United Kingdom in October 2013, the Essex Police introduced a new operation that, over a 19-month period, targeted a total of 6,000 200-meter-radius areas. Every week a different set of areas was chosen, with each area receiving an average of 10 additional minutes of police presence per day. Together, these areas represented the locations where crime (specifically, private-dwelling burglary) occurred in Essex during this period.

The policy had two features that make it uniquely valuable from an econometric perspective. Firstly, the weekly choice of targeted areas was determined by a simple and rigid rule. We can use this rule to identify areas that would have been targeted in the period immediately preceding October 2013, had the policy already been in place. Secondly, the time variation of the increase in policing for the selected areas was discontinuous and driven by inflexible organizational constraints: it started on a Friday and lasted exactly seven days. The policy allows us to convincingly isolate exogenous variation in police presence, account flexibly for the possibility of correlated shocks in the crime propensity of an area, and construct various counterfactuals that are not influenced by the policy.

We fail to find a decrease in crime that corresponds to the increase in police patrolling induced by the policy. Furthermore, our estimates are precise enough to be able to reject an elasticity half the size of those found by Di Tella and Schargrodsky and by Draca, Machin, and Witt.

In light of these findings, it is particularly important to understand whether our result might be caused by deficiencies in the estimating sample or in the measurement of the crime and police deployment variables. We argue that this is unlikely, for three reasons. First, compliance with the policy was very high. Second, the large temporal and geographical breadth of the policy provides sufficient statistical power to rule out relatively small effects. Third, our research is unusual in that it measures police presence through a GPS-based dataset that records the location of each police officer every 2.18 minutes. We argue that the precise measurement of the police presence variable makes measurement error an unlikely explanation for our findings.

We explore other potential explanations for our main finding, and end up rejecting all of them. First, we find that the additional police deployment created by the policy occurred during the hours of the day when crime typically takes place. This suggests that temporal misallocation of patrolling intensities is not a likely explanation for the zero finding. Second, our main finding is robust to varying the type and circumstances of the crime. For instance, we find the same effect when focusing only on crimes cataloged to have taken place on the street (where patrolling should, in principle, be more effective). The effect is also the same regardless of the preexisting crime or patrolling intensities of an area.

It would be surprising if, even under the most favorable circumstances, police street deployment was completely unable to deter crime. More interesting from a policy perspective is the question of whether police patrolling as typically practiced provides a sizable contribution to crime prevention. We believe that the natural experiment that we exploit generates evidence that is more easily extrapolated to the circumstances and levels of police deployment that are the norm. The evidence is consistent with the elasticity between police and crime measured at the daily level being relatively small.

Vigorous debates are taking place in many jurisdictions regarding the optimal size of police budgets, as well as the optimal allocation of resources inside departments. Concentrating resources in a small number of locations or incapacitating criminals by responding rapidly to calls for service have been argued to be effective strategies in combating crime, and may compete in terms of resources with the policy of using general-purpose patrolling to thinly spread police presence through large areas. While a cost-benefit analysis of alternative uses of police resources is beyond the scope of our research, we contribute to this debate by providing estimates of the relationship between police and crime that may be more generalizable, and therefore more useful, to the policymaker than those in the existing literature.

NOTE: