SCIENTIFIC COMMUNICATION

Targeting Knife-Enabled Homicides for Preventive Policing: A Stratified Resource Allocation Model

Vincent Harinam [1] · Lawrence W. Sherman [1]

Abstract

Research Question How can police translate differing risk levels for knife homicide into a resource allocation model that follows the evidence?

Data The data for this publication are taken from the open access tables published in this journal by Massey et al. *Cambridge Journal of Evidence-Based Policing*, 3:1-20, (2019). Those data show the linear relationship between the number of non-fatal knife assaults in a lower super output area (LSOA) in 1 year and the risk of a knife-enable homicide in the subsequent year, as well as how many of the 4835 LSOAs fell into each of five levels of increasing homicide risk.

Methods The data from Massey's research are re-calculated to show how a hypothetical number of 15-min police patrols could be allocated across all areas on the basis of a *combination* of knife-enabled (KE) homicide risk level and the volume of LSOAs at each of the five levels of knife homicide risk. We display these results using both tables and multi-layered "wedding cake" images to show the size of different dimensions of each level, including proportion of total homicides and directed patrol frequency per LSOA at each of the five risk levels.

Findings Based on the hypothetical allocation of 10,000 patrol visits of 15 min in length, the highest risk group, with a forecasted 6% of all KE homicides, would receive 600 police patrols, divided by the 41 LSOAs at that risk level = 15 patrols across every 10 days. At the lowest level of risk, the 2787 LSOAs would share the 3000 patrols that a group of LSOAs would recieve for having 30% of homicides, which equals 1.1 patrols every 10 days. The hypothetical premise is that every LSOA gets some patrol, but the highest risk areas get 15 times more patrol to follow the evidence of risk. The formula is to (1) allocate resources by proportion of homicide at each risk level; (2) divide the allocated resources by the number of areas in each risk level group; and (3) allocate the resulting resources per day to each area in each of the 5 levels.

✉ Vincent Harinam
vh315@cam.ac.uk

[1] Institute of Criminology, University of Cambridge, Cambridge UK and Centre for Evidence-Based Policing Ltd., Cambridge, UK

Conclusion Police face difficult tradeoffs between targeting more policing to *fewer* areas of higher risk (with more efficiency) or to *more* areas of lower risk (with more effectiveness). The use of a formula combining risk and volume can help guide such decisions, illustrated by a layered "wedding cake" visualization for gaining clarity and legitimacy in communications.

Keywords Preventive policing · Homicide · Weapons · "Wedding cake" graphics

Introduction

The discovery of spatial concentrations of most serious crimes in a small number of "hot spots" (Sherman et al. 1989) as a "law of crime concentration" (Weisburd 2015) has provided a widespread targeting strategy for evidence-based policing. The effectiveness of these strategies, using both preventive patrol and problem-oriented policing, has been tested in a wide range of randomized trials and quasi-experiments, as systematically reviewed by Braga et al. 2012, 2019). The most recent systematic review found that 62 of the 78 reported tests of hot spots policing demonstrated crime reductions from increased police activity, compared with no increases in similarly "hot" locations (Braga et al. 2019). While the average reduction in crime volume has been relatively modest, the large number of crimes in those locations appears to make such investments cost-effective (e.g., Ariel et al. 2016).

The testing of police activity at *all* targeted hot spots (above a fixed threshold) has grown substantially, but the same cannot be said of research on targeting high-crime locations with *differing levels* of risk and resource allocation. Despite the consistency with which hot spots remain on the rank-ordered list of "hottest" locations from year to year (Weinborn 2017), there has been little research on the extent to which policing resource levels could *vary across* hot spots during any given year. A binary model of "hots" and "not hots" for police resources has the appeal of simplicity. But it may be too simple when policing rare but serious events, such as murders and serious violence, in which a substantial part of the problem is spread across an entire city without concentrations (as in London). In other words, hot spots policing has been treated as a *binary* strategy for resource allocation, and not as a *stratified* layer strategy in which several levels of risk can receive proportionately different levels of policing.

The issue of what proportion of serious, or all, crime in a city is even concentrated into the identifiable and patrollable hot spots is as old as the evidence on this strategy. In the first randomized trial of hot spots patrols, for example, Sherman and Weisburd (1995: 634) selected a sample of 110 of the hottest hot spots in the city, which collectively accounted for only 10% of all of the calls for service about crime in Minneapolis over 1 year. Thus, even if the extra patrols had eliminated all crime at the 55 hot spots randomly selected to receive the extra patrol, it would still have only reduced crime calls by a maximum of about 5%. Few other studies of hot spots policing have even estimated the maximum contribution of a hot spots strategy to an overall crime reduction strategy.

This problem is even more important for the crimes in public places that attract the greatest levels of public concern and media attention: the high-harm, low volume crimes of murder or potentially fatal assaults. Massey et al. (2019), for example, reported that 31% of the knife-enabled (KE) murders in London in a recent year occurred within the 2787 small areas which had not had a single stabbing in the prior

year. Those areas were definitely not hot spots of knife violence. But they did suffer a major proportion of a crime of great importance. To the extent that extra police activity is targeted only at locations with the highest level of risk for high-harm crime, there could be a cause for concern about withholding extra policing from a majority of areas in which a third of such crimes occur.

The question of how narrowly or broadly to allocate knife crime prevention resources across London was addressed by Massey et al. (2019: Figures 7 & 8), using pie charts to attempt to illustrate the choices that could be made from the analysis. Whether these graphics help police leaders to interpret the findings for making operational decisions is unclear. Yet some feedback from Cambridge classes for police leaders suggests that the pie charts are not enough to guide precise resource allocation decisions across 4835 areas. Nor is it enough to simply have the evidence to identify five risk levels of progressively higher risk of homicides in some areas that others.

In this scientific communication, we attempt to provide further guidance on resource allocation based on risk stratification that guides a proportional allocation of resources. Our approach is based on the Massey et al. (2019) demonstration that the level of risk of homicide is clearly, and steeply, stratified across the Metropolitan Police area. Our aim is to apply the London data in a resource allocation formula that combines the risk level of each LSOA with the volume of LSOAs at each of the five risk levels. Our hope is that the allocation of patrol presence to areas of stratified risk can be accomplished with increasing precision, with maximum benefits for the community in reduced harm and violence.

Research Question

Our research question is this:

How could the evidence of widely varying homicide rates across five strata of risk levels among 4835 areas best used to decide how much patrol or other preventive activity each area should receive?

We frame the question as a hypothetical (simulation) analysis in which we assume that across London each day at least 1000 patrol car visits of 15-min duration can be delivered for the purpose of preventing knife violence and homicide. The round number of 1000 is based on clarity of illustration, rather than any professional judgment about how best to use Metropolitan Police Service (MPS) resources.

A patrol visit length of 15 min, in contrast, is based on substantial empirical evidence that 15-min visits create more lasting crime reduction effects than shorter visits of police cars to hot spots. Koper (1995) first reported this finding in his "Koper curve," showing that the longer the police stayed in a Minneapolis hot spot, the longer the hot spot remained free of crime or disorder after police left, up to the end of a 30 min follow-up period. Williams and Coupe (2017) showed in a randomized trial in Birmingham UK that three daily foot patrol visits to a hot spot lasting 15 min each (totaling 45 min) produced less crime in the same hot spots than nine daily visits of some 5 min in length (also totaling 45 min). Ariel et al. (2020) reported that in a randomized trial in the London Underground, 15-min patrols on over 50 randomly selected hot spot platforms 4 days a week reduced crime across 7 days per week, compared with having no regular patrols on some 50 other equally "hot" high-crime

frequency platforms. Most recently, Barnes et al. (2020) have reported that the effects of one randomly assigned 15-min visit produced residual deterrent effects lasting for up to 4 days of no further patrols.

Shaping the research question on a hypothetical allocation of 15-min patrols, even at the rate of less than one patrol per day, provides an evidence-based simulation linked to substantial prior research. Answering that question requires as much "translation" as possible for such important basic discoveries as the Massey et al. (2019) research.

Data

In 2019, Detective Chief Inspector John Massey of the Metropolitan Police sought to understand how accurately recorded locations of knife crimes in 1 year could be used to forecast fatal stabbings in the following year (Massey et al. 2019). At the level of lower super output areas (LSOA), Massey collected all recorded non-fatal knife-enabled assault in London in the financial year 2016–2017. He then collected all knife-enabled homicides in the following financial year (2017–2018). With these geospatial figures, Massey calculated the mean probability of a fatal stabbing in 2017–2018 for each level of frequency of knife offenses in 2016–2017 within each group of LSOAs. This research found that the more knife-related (but non-fatal) *assaults* an LSOA had in year 1, the higher the probability it had of having a knife homicide in year 2.

Our analysis of the Massey et al. (2019) data is comprised of three information elements that we use to stratify risk for a resource allocation formula that can be derived from the evidence:

1) The raw number of LSOAs in each category of year one knife assault frequency, from level I to V
2) The probability of a knife-enabled homicide within each of the LSOAs in each of the five risk levels for KE homicide
3) The percentage of total knife-enable (KE) homicides city-wide that occurred in each of the five levels

To these ingredients, we add the assumption that *on each day* only 1000 patrols (lasting 15 min in each LSOA) can be allocated to preventing knife violence. We then present a model for allocation those patrols over a 10-day period (10,000 patrols), in order to assign patrol frequency of less than one visit per day to the lowest risk LSOAs in the MPS area.

We also display these results using a "wedding cake" graphic that the first author developed by manual construction using Microsoft Word. While no bespoke program exists to automatically construct such graphics, they can be replicated without an expensive program or a high level of technical expertise. The main tool is the "basic shapes" and "flow chart" features in the "shapes" button on the insert ribbon in Microsoft Word.

The key to designing this visualization is proportionality. The graphic is most informative when the physical size of a layer or the width of a cone corresponds, at least roughly, to the value of the data point that is attributed to it. Thus, the lowest layer of the cake must be physically larger than the others when it represents the targeting model that covers the most spatial units. When the graphic illustrates the differences in

Table 1 Five levels of risk of homicide by percentage of homicide at each level

Risk level numeral	N of LSOAs	N of knife homicides	KE homicide risk per LSOA	Percent of All KE homicide
V	41	6	15%	6%
IV	35	3	9%	3%
III	714	30	4%	31%
II	1258	28	2%	29%
I	2787	30	1%	31%
Total	4835	97	(N.A.)	100%

(Source: Massey et al. 2019)

patrol frequency recommended for highest versus lowest risk levels, the opposite proportionality is required—showing that highest risk LSOAs get more frequent visits per LSOA than lower risk LSOAs.

Findings

Table 1 shows (from Massey et al. 2019) that there is a relatively small number of LSOAs at the highest risk of homicide and a majority of all LSOAs at the lowest risk of homicide. But it also shows that targeting only the highest risk LSOAs would be unlikely to prevent much homicide. What is different in Table 1 here from Table 2 in Massey et al. (2019) is that here we have *separated the risk level* for each of the five levels rather than aggregating it as the targeting adds more LSOAs. In the 2019 table, the risks for each lower level were aggregated with the high-level risks. That approach has many uses in understanding the problem. This table, in contrast, estimates the separate risks at each risk level independently from any aggregations with the other risk level.

The new findings in Table 1 demonstrate that there are different numbers of LSOAs at each risk level and that the volume of LSOAs by risk level can be used to allocate a bespoke level of resources for that risk level. By dis-aggregating the homicides in each risk level, the table allows a formula to be developed for resource allocation that is based on the total homicides at each risk level. Level IV, for example, not only has a substantially lower risk level than level V, it also has fewer LSOAs than are found at level V.

By combining the number of LSOAs in each of the five risk level groups with the raw numbers of LSOAs in each group, we find the ingredients for allocating resources by LSOA by risk level.

Table 2 shows how we can compute the percentage allocation of policing resources based on the percentage distribution of homicides in each LSOA risk group, combining the homicide risk per LSOA with the total proportion of all homicides. By hypothetically allocating 10,000 patrol visits across all 4835 LSOAs, we report the application of the formula used in Table 2. Our metric of how many patrol visits should be made every *10* days is set at that level to minimize the use of decimal points. At the lowest risk level, one patrol every 10 days is allocated, rather than the (potentially

Table 2 Five levels of resource allocation per LSOA based on risk and total share

Risk level numeral	Percent of KE homicide	Total visits per 10 days	N of LSOAs	Visits per 10 days per LSOA
V	6%	600	41	14.6
IV	3%	300	35	8.6
III	31%	3100	714	4.3
II	29%	2900	1258	2.3
I	31%	3100	2787	1.1
Total	100%	10,000	4835	(N.A.)

meaningless, or at least confusing) statement of one-tenth of a patrol per day. If patrol times should be based on empirical evidence, the full 15 min appears necessary whenever it is provided and with whatever frequency or infrequency.

The formula by which the farthest to the right column is derived is as follows:

1) Allocate resources to each group by proportion of *homicide* at each risk level.
2) Divide the allocated resources by the number of *areas* in each risk level group.
3) Allocate the resulting *patrol visits per 10 days per area* to each area in each of the 5 levels.

Visualizing Risk and Resource Allocation

In Fig. 1, the risk levels per LSOA drive the five layers of the wedding cake from the lowest to highest risk. The wedding cake image clearly shows the largest number of LSOAs to be at the lowest level of risk, as indicated on the far right-hand gradient. Over half of all 4835 LSOAs are at a 1% risk level of homicide in the year being forecast, while under 1% of the LSOAs are found in level V at the level of a 15% chance of having a KE
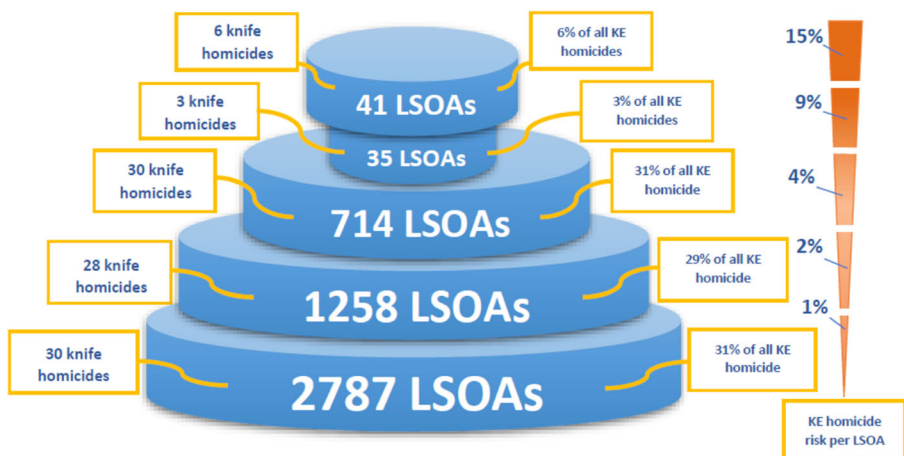

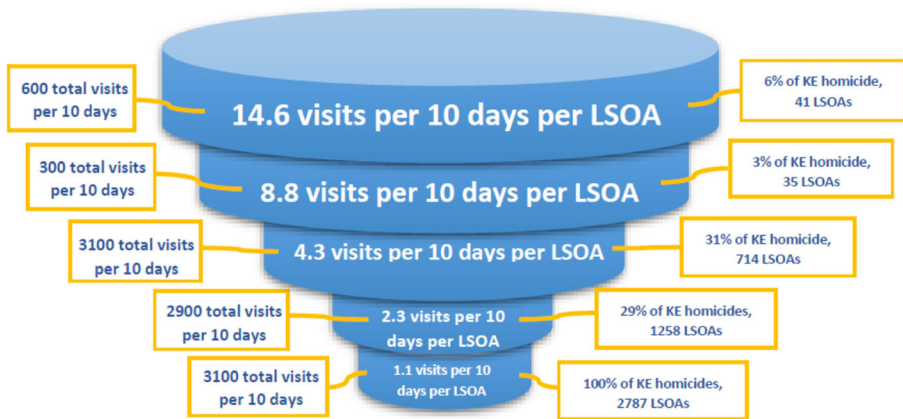**Fig. 1** Risk levels and homicide proportions in London 2017

Fig. 2 Resource allocation by risk and proportion of homicides: a hypothetical distribution

homicide in that year. It provides visual evidence to suggest that if the risk is 15 times higher in level V LSOAs than in level I LSOAs, then perhaps the LSOAs at level V should be given 15 times as much policing against knife crime as those at level I.

Figure 2 displays the results from Table 2, which applies the formula that combines risk levels with numbers of units where each risk level applies. Its hypothetical premise is that 1000 police visits of 15 min can be delivered by some patrol unit each day in London. This premise is plausible because even at 4 such visits per day per patrol unit, 250 patrol units per day delivering 4 visits spread across three shifts, or 83 patrol visits per 8-h shift, can be deliver the allocated level with existing levels of patrol car coverage.

Figure 2 shows a paradox: that half of all patrol visits would be delivered to the two lowest risk levels (I and II), featuring only 1% or 2% risks of homicide. That allocation is based on the fact that those levels actually suffer half of all homicide. To do otherwise would constitute a decision to focus on only the other half of all KE homicide, which is found in the higher risk areas (levels III through V).

The more compelling visual impression in Fig. 2, however, is that the amount of policing per LSOA rises in direct proportion to the risk levels of those LSOAs. That distribution creates an "upside-down" cake, by which the higher the risk level, the more policing *per LSOA* is allocated. This may be a basis for changing the units for such patrols from generalist to specialists and allow a specialist unit to stay close to high-risk areas for most of a shift. However, the differential policing is assigned; it is more clearly based on a transparent risk stratification than previous patrol allocation models.

Conclusion

This communication demonstrates two things. One is a formula for calculating resource allocations to prevent rare but high-harm crimes. The other is a method for visualizing both risk stratification and resource allocation.

The formula suggested in this paper is certainly not the only possible way to allocate the available resources. Other models with other assumptions are certainly possible. Other outcomes besides knife violence, for example, could be woven into a risk analysis: robbery, rape, affray, and even vehicle theft might all be deterred by police

visits. If the outcomes are linked to a single resource allocation, then that allocation could be based on that broader range of outcomes. Alternatively, the formula could discount the harm levels in lower risk areas, and raise their value in higher risk areas, for reasons that may be important in a larger policing picture: police legitimacy may require more (or less) policing in some areas, for example, independent of the pure risk levels. The main point of reporting this formula is to show how any formula might work by applying evidence systematically.

That same point would apply even if the decision was to target all patrol to LSOAs above a certain threshold of risk, such as a 5% chance of a homicide. That would only target 9% of total homicide, but it would increase available patrols for those areas by tenfold or 9100 patrols per 10 days. That choice is impossible to make solely based on current evidence, but future studies could show that extra patrols only prevent homicide in the highest risk areas—a finding echoing Ratcliffe et al. (2011) that showed foot patrol only prevented violence in Philadelphia above a threshold level of frequency within patrol beats.

The "wedding cake" models for visualizing stratified resource allocation may have far broader implications for police strategies, both within and beyond the domain of geographic targeting of resources. Domestic violence cases, for example, are steeply stratified by harm levels, with a vast base of events that require no medical attention, rising to a tiny proportion of all such events that result in death, rape, or torture. The allocation of police resources in that domain is far from proportionate to the harm levels, with steep ratios of under-investment at the highest harm levels.

Using this kind of graphic avoids the constraints of more traditional shapes, such as pyramids, which do (or may) not match the shapes of the distributions. A picture may be worth a thousand words. And a wedding cake may be worth far more than a set of tables that display numbers that create no picture.

References

Ariel, B., Weinborn, C., & Sherman, L. W. (2016). "Soft" policing at hot spots—Do police community support officers work? A randomized controlled trial. *Journal of Experimental Criminology, 12*(3), 277–317.

Ariel, B., Sherman, L. W., & Newton, M. (2020). Testing hot-spots police patrols against no-treatment controls: Temporal and spatial deterrence effects in the London underground experiment. *Criminology, 58*(1), 101–128.

Barnes, G. C., Williams, S., Sherman, L. W., Parmar, J., House, P., & Brown, S. A. (2020). Sweet spots of residual deterrence: A randomized crossover experiment in minimalist police patrol.

Braga, A., Papachristos, A., & Hureau, D. (2012). Hot spots policing effects on crime. *Campbell Systematic Reviews, 8*(1), 1–96.

Braga, A. A., Turchan, B., Papachristos, A. V., & Hureau, D. M. (2019). Hot spots policing of small geographic areas effects on crime. *Campbell Systematic Reviews, 15*(3), e1046.

Koper, C. S. (1995). Just enough police presence: Reducing crime and disorderly behavior by optimizing patrol time in crime hot spots. *Justice Quarterly, 12*(4), 649–672.

Massey, J., Sherman, L. W., & Coupe, T. (2019). Forecasting knife homicide risk from prior knife assaults in 4835 local areas of London, 2016–2018. *Cambridge Journal of Evidence-Based Policing, 3*(1–2), 1–20.

Ratcliffe, J. H., Taniguchi, T., Groff, E. R., & Wood, J. D. (2011). The Philadelphia foot patrol experiment: A randomized controlled trial of police patrol effectiveness in violent crime hotspots. *Criminology, 49*(3), 795–831.

Sherman, L. W., & Weisburd, D. (1995). General deterrent effects of police patrol in crime "hot spots": A randomized, controlled trial. *Justice Quarterly, 12*(4), 625–648.

Sherman, L. W., Gartin, P. R., & Buerger, M. E. (1989). Hot spots of predatory crime: Routine activities and the criminology of place. *Criminology, 27*(1), 27–56.

Weinborn, C. (2017) From hotspots to harmspots, an empirical exploration of crime concentrations and crime harm in the United Kingdom. Unpublished Ph.D. dissertation, Institute of Criminology, University of Cambridge.

Weisburd, D. (2015). The law of crime concentration and the criminology of place. *Criminology, 53*(2), 133–157.

Williams, S., & Coupe, T. (2017). Frequency vs. length of hot spots patrols: A randomised controlled trial. *Cambridge Journal of Evidence-Based Policing, 1*(1), 5–21.

Vincent Harinam - PhD candidate, Institute of Criminology, Univeristy of Cambridge.

Lawrence W. Sherman - Project Manager, Cambridge Centre for Evidence-Based Policing.