

Patterns in the supply and demand of urban policing at the street segment level

Toby Davies^a and Kate Bowers^a

^aDepartment of Security and Crime Science, University College London, 35 Tavistock Square, London, WC1H 9EZ, UK.

ARTICLE HISTORY

Compiled March 27, 2019

ABSTRACT

Policing plays a critical role in crime prevention, incorporating both deterrent and responsive activities. Since many policing activities require the physical presence of officers, a crucial issue for police effectiveness concerns the extent to which officers are located where they are needed. Operationally, this can be framed as a ‘supply and demand’ issue, where the aim is to match the supply of police resource to the demand for service. In this paper, we examine this issue for a 5-month period in London, UK, using police vehicle tracking data and call-for-service records. We examine the extent to which supply and demand are aligned at the street segment level, and build a statistical model which seeks to explain the disparity between the two quantities in terms of network structure. We find that police activity is distributed unevenly, with over-supply on more central streets, and discuss the implications of this for police practice.

KEYWORDS

crime; policing; street networks; officer tracking

1. Introduction

Policing plays a central role in modern society, as one of the foremost means by which the safety and wellbeing of the public is maintained. Traditionally, its primary function has concerned the prevention and detection of crime, incorporating both responsive and deterrent activities. In many modern policing models, these are augmented with additional responsibilities in other areas, with officers playing a more prominent role in responding to non-criminal incidents and providing reassurance or assistance to the public (see ?). Although these activities are quite diverse, what is common to all is that they involve the physical presence of police resources, either in order to respond to incidents or to act as a visible authority.

The fact that so much of policing depends on physical presence implies that the location of officers is a crucial determinant of both effectiveness and efficiency. Put simply, a force will be more efficient if its resources are located in places where they will have greater impact (on crime or any other aspect of performance). This can be framed as an issue of supply and demand: the objective is to match spatially the ‘supply’ of policing (*i.e.* officer time) with the ‘demand’ for it (*i.e.* requirement

for service). In practice, police command and control seeks to achieve this through both long-term planning (of patrol routes and priority areas) and real-time dispatch management. In many cases, this is also supported by analytical work or ‘predictive policing’ systems, which seek to identify potential crime locations prospectively (see ??).

Despite the clear importance of spatial resource allocation, however, relatively little is known about the extent to which police activities are aligned with the demand for them in practice. The most prominent reason for this has traditionally been the paucity of data concerning the locations and movements of police officers. Previous empirical work has been based on either small samples of data or records of intended (rather than observed) activities. Otherwise, they have been undertaken at aggregated levels which, whilst informative, have masked interesting variation at the micro-level (see, for example, ?). However, the rapid proliferation within policing of GPS-enabled mobile devices means that this barrier has largely been removed: the locations of police units are now systematically monitored and recorded with high spatio-temporal accuracy. This means that it is possible to analyse the micro-level spatial characteristics of police activity at a large scale, and, furthermore, to compare its distribution against that of policing demand.

Measuring both need and provision in any societal context can be challenging. In policing terms the ‘supply’ element can take many forms. For example, foot patrol provided by officers on the beat is a different type of provision to the potential reassurance or deterrence offered by a marked police vehicle on its way to respond to an incident. This differs again from the provision of service via the detection of a perpetrator or direct communication with a victim. Because of this diversity, previous research focusing on policing supply has tended to focus on one type or the other, or assessed such activities distinctly (*e.g.* ???). Further, the extent, timing and duration of policing that is necessary for a function to be successfully performed will vary across activities. Hence the ‘dosage’ of policing that is necessary to break up a Saturday night city centre fight is likely to differ from that necessary to deter a lone-acting burglar from breaking into a residential home during the daytime on a week day. There is further evidence that, from a prevention point of view, anti-crime measures are unlikely to offer linear gains; that is, there is not necessarily a proportional increase in reductive effect as more resources are invested (?). The concept of dosage is therefore complex, and we return to this issue later.

In this paper, we examine the correspondence between policing supply and demand at the street segment level for the city of London, UK. This is done using a dataset which provides GPS location records for all police vehicles over a 5 month period in 2015. By processing this, we are able to measure accurately the volume of time spent by police vehicles on each street segment in the city. This is then compared against the volume of demand generated on each segment, measured in terms of calls for service: these represent all incidents which require a physical police response. We examine the association between these values for each street segment, and identify patterns in the discrepancy between the two. To our knowledge, this is the first study which examines police activity at this level of granularity, and which considers the extent to which routine police activity conforms to operational demand.

The aim of this paper is exploratory, in the sense that we seek only to identify patterns in the provision of policing activity, rather than to compare them to any form of idealised situation. Although a naive assessment might suggest that optimal police behaviour would result in perfect alignment between supply and demand, this is clearly unrealistic in practice: officers engage in a diverse range of activities, and calls

for service are not the only considerations that govern their movement. Furthermore, any purposeful movement (*e.g.* in response to a call) necessarily involves travel along a route, meaning that some places are visited simply because they lie on a path between other locations. This is complicated further by the fact that all movement takes place in a complex urban environment - governed largely by the street network - which shapes not only patterns of travel but can itself influence the distribution of demand. Our analysis also examines the influence of this aspect of urban form on the balance between supply and demand.

Although we do not approach our analysis with any particular notion of optimal or desirable behaviour in mind, exploration of this issue does have significant potential consequences for real-world policing. Locations in which supply and demand are misaligned - in either direction - represent possible causes for concern, and spatial patterns among these may be indicative of systematic inefficiencies in police allocation strategies. These could be remedied by introducing alternative behavioural principles within policing; that is, by making changes to the way in which officers move around cities. It may be possible, for example, to suggest route-finding heuristics which give rise to more appropriate coverage patterns, or to identify optimal locations for officers to wait between calls. These could be implemented using intelligent deployment systems, which would be well-aligned with the increasing use of data-driven software systems within policing. Our focus in the paper is explicitly operational: we are concerned with the efficient implementation of policing, rather than its ecological grounding or potential feedback effects with crime. Although the demand for policing is a complex phenomenon, the fundamental challenge facing police agencies is to match their activities to the demand as they observe, and it is this aim with which we are concerned.

We will begin the paper by providing a brief review of previous research concerned with police resource allocation and officer tracking. We will then motivate the latter part of our analysis by outlining evidence concerning the role of urban form in shaping human activities. Following this, we will describe the technical details of our work, including the data we use and the way in which it is processed. Finally, we will show empirical results, in which we will identify spatial patterns in the relationship between supply and demand, with particular emphasis on the role of the street network.

1.1. *Police movement and its tracking*

Throughout the history of modern policing, the way in which the police exert their physical presence, and the way it is measured, has gone through a number of stages. The reasons for this are both technological and political: logistical advances have provided greater capabilities, while the intended function of physical patrol has evolved across a number of operational models. These changes have been documented extensively (*e.g.* ??), and the strengths and weaknesses of the various approaches discussed. It is clear, however, that much research in this area is broadly qualitative: compared with the ‘outcomes’ of policing (*e.g.* crime or calls for service), police behaviour itself has received little quantitative attention. In this section, we will give a brief history of police resource tracking.

A central component of the modern model of policing introduced by Peel in 1829 was the notion that the police should act as a visible representation of authority (see ?). The majority of their presence therefore came in the form of preventative patrol, designed to deter crime by providing guardianship with the ability to interrupt and punish wrong-doing. In keeping with the state of technology at the time, these

patrols were organised quite simplistically: policing areas were divided into small areas known as ‘unit beats’, each of which was the responsibility of a single officer. Patrols were then done on the basis of a ‘fixed point’ system: the police maintained a set of static locations (eventually equipped with telephones), which officers would travel between in some order, ‘checking in’ with their base station at each one. Although such a system had numerous shortcomings, it was extremely expedient with respect to resource tracking: officers were constrained to small areas, and their positions were regularly known with certainty.

This model of policing persisted for some time, but eventually responded to technological developments, most notably the introduction of motor vehicles and police radios (see ?, for a summary of the evolution of police organisation in the twentieth century). These afforded officers much greater freedom and agility, and brought clear benefits in terms of their versatility and ability to respond quickly to incidents. Individual units typically operated over much larger areas, and activity could be coordinated much more readily. Inevitably, however, this affected the model of policing: the fact that officers *could* respond to incidents much faster meant that the emphasis shifted towards reactive activity. Officers therefore occupied a ‘crime fighting’ role, spending a greater proportion of time ‘hopping’ between calls rather than engaging in planned preventative activity. Additionally, officers were given much greater autonomy, with discretion to patrol as they saw fit between calls. This activity was largely unmonitored: aside from being in touch via radio, relatively little was known about officer movements.

Over time, it became recognised that the general lack of accountability in resourcing practices was unacceptable, and police forces began to adopt a more business-like approach. As resources became more rigorously tracked, and budgets increasingly constrained, efficient deployment became an imperative, and forces began to explore how to ensure that their resources were better aligned with demand. In the UK, two documents - an Audit Commission report (?) and the ‘Tackling Patrol Effectively’ guidelines (?) - sought to re-establish the role of patrol in policing, highlighting the need to tailor activity to demand and to restore the balance between reactive and preventative activity. In addition, a ‘Value for Money’ report (?) advocated formal auditing, leading to the introduction of ‘Activity Based Costing’ surveys, via which officers were required to log their activity. While theoretically providing comprehensive records, however, the form of these - written logs to be completed every 15 minutes - rendered them bureaucratic and vulnerable to manipulation, and they were soon abolished.

A further development within policing in recent years concerns the increased role played by analysis, both of crime and of police performance. The use of Geographical Information Systems (GIS) has become commonplace, and provides a means to gain a rigorous understanding of exactly where and when crime happens. Such analysis has revealed striking regularities in the distribution of crime, the most significant of which concern the fact that crime is highly concentrated in space and time: a small proportion of locations account for the majority of incidents (?). This has inspired a number of widely-used policing strategies which seek to concentrate crime prevention effort in these areas, motivated by evidence which suggests that police presence in such locations has a deterrent effect on crime (???). This has brought about a renewed focus on planned, pro-active activity, and the desire to place ‘cops on the dots’ in this way is a central principle of the performance-driven management regime COMPSTAT (?) and its variants.

Despite the perceived importance of patrol, however, there has still been relatively

little monitoring of where police officers actually spend their time in practice; a key component of the COMPSTAT ‘equation’. The reason for this has traditionally been primarily technical: no means existed to track activity in a systematic and unobtrusive way. These technological barriers have, however, been removed as a result of the growing proliferation of GPS-enabled mobile devices throughout policing. Nevertheless, few studies have utilised this data in the examination of police behaviour. Those which have focus primarily on targeted interventions, using movement data to measure the ‘dosage’ applied to particular locations (*e.g.* ??). In addition, studies have also used such data to identify behavioural patterns in the movement of officers, such as their routing choices (?) or locational preferences (?). There have, however, been no attempts to examine patterns in the physical presence of officers across the entire business of the police. For those concerned with improving performance across all dimensions of police tasking, however, this is a crucial issue.

1.2. *The influence of urban form*

In simple terms, the movements of the police officers are governed by two factors: the locations to which they go, and the routes that they take between them. With respect to both of these, urban form - and, in particular, the street network - is likely to play a role in shaping behaviour. The street network is one of the primary determinants of the configuration of cities: it is the substrate around which buildings and facilities are arranged, and it is the structure on which people move. As such, it provides a means to understand the organisation of human activity, and to characterise particular places. Since streets vary substantially in terms of where they are and how they are used, it is to be expected that relationships may be observed between their properties and the issues of interest in this work.

The structure of street networks has attracted considerable research attention in recent years, and a number of studies have examined the relationship between network properties and human activities. One of the central premises of such work is that some streets can be considered to be more ‘central’ than others (where centrality is defined in a number of senses), and that this influences how they are used in routine urban activities. By representing the street network formally as a graph, these notions of centrality can be quantified using explicitly-defined metrics. The properties of real-world networks can then be compared with the spatial distribution of human activity, as demonstrated in two studies which measure the association between street centrality and economic activity in Barcelona (?) and Bologna (?). In each of these cases, positive relationships were observed between the locations of retail services and an array of centrality metrics: that is, economic activity tends to be located on more central roads. This accords with research based on the closely-related ‘space syntax’ approach, which has explored the relationship between urban form and social dynamics in depth (??).

There are strong theoretical and empirical reasons to anticipate that similar relationships can be observed for crime. On one hand, that such relationships should exist is predicted by a number of theories within environmental criminology, most notably crime pattern theory (?). This suggests, among other things, that the distribution of crime can be rationalised in terms of the ‘awareness spaces’ of potential offenders: that places with which people are more familiar are more likely to be considered as locations in which to commit crime. The extent to which places fall within the awareness spaces of citizens is, of course, influenced by their position in the street network. There is likely to be greater awareness, for example, of locations which are used more

regularly (perhaps because they lie on popular routes between places) or are more accessible. On the basis of this argument, therefore, it is to be expected that - with all else equal - certain streets will be at greater risk of crime by virtue of their position in the network.

A number of studies have explored the relationship between crime and the street network, using analytical methods of varying sophistication. After several early studies which considered connectivity at the area level (*e.g.* ??), more recent work has focussed on the street segment as the unit of analysis. ?, for example, examined the relationship between crime risk in Ridge Meadows, Canada, and two segment-level properties. They found that both accessibility (measured as the number of neighbouring segments) and flow (measured using road classifications) were positively associated with risk, after controlling for other variables.

Crime risk has also been examined using the ‘space syntax’ approach, one of the central principles of which is that the character of street segments can only properly be understood in the context of the rest of the network and their position in it. Applying the approach to crime data from London, UK, and an urban area in Australia, it was found that crime was positively associated with connectivity, but negatively associated with a metric referred to as ‘integration’ (?). In addition, particular consideration was given to *cul-de-sacs*, which represent an extreme case as a result of their isolation. In ?’s analysis, secluded segments such as these were found to be at particularly high risk, although it was emphasised that their physical geometry must also be taken into account. This is a finding which runs contrary to pattern theory, as well as to empirical results found at the individual property level by ?.

Some years later, ? explored these issues further using data from Merseyside, UK. Although the independent variables considered - number of connections and administrative classification - are similar to previous studies, their approach is distinguished by the use of a statistical framework which accounts for hierarchical spatial nesting in the data. After controlling for other variables, it was found that both connectivity and administrative classification were positively associated with burglary risk, with *cul-de-sacs* experiencing particularly few offences. This approach was extended further by ?, who introduced a more formal quantitative framework for the measurement of network properties. This facilitated the use of more sophisticated network metrics, which correspond more closely to underlying behavioural theories. One of these - *betweenness* - which measures how frequently individual segments will be used in travel through the network, was found to be positively associated with burglary risk. These results, and those of the majority of other studies mentioned, suggest that crime tends to occur on better connected and more heavily-used segments. It follows then that the demand for policing (*i.e.* calls for service) at the segment level is likely to vary with network properties such as these.

1.3. Route choice in urban environments

While the relationship between the street network and crime has only recently been rigorously examined, its role in human movement more generally is a long-standing and mature field of research. Transport is, of course, the primary function of street networks, and it is natural to expect that its structure should influence patterns of both pedestrian and vehicular travel. In much the same manner as invoked by crime pattern theory, the overall patterns of movement in a city are formed through the accumulation of individual journeys, each of which involves the selection of a path

through the network. The study of route choice (for example of offenders or of police vehicles) seeks to establish which factors influence this selection, many of which can be expressed in terms of network properties.

Traditional approaches to route choice modelling have been predicated on the principle that individuals seek to optimise their travel time when moving between places. At its simplest, this implies that they will take the shortest possible path, but this approach can be refined by taking into account factors such as congestion (as in classic traffic assignment models; see ??). The assumption of optimal behaviour, however, has been challenged by other authors (?), who suggest that it fails to reflect the complexity of the route choice process. In response, more individual-centric models have been introduced, where the choice between alternative routes is influenced by an array of factors, including subjective perceptions of routes (?) and spatial ability (?).

An alternative perspective has been inspired by findings from spatial cognition research, which suggest that particular features of urban space play a key role in navigation (??). These have given rise to ‘anchor-based’ models of route choice, in which journeys are framed in terms of transitions between landmarks or other prominent locations (?). Models based on these principles have been shown to represent real-world behaviour more accurately than time-optimising methods (see ?), and a range of similar behavioural principles have been explored. Although the details of these are beyond the scope of this paper, what is clear among all of these approaches is that movement behaviour is highly constrained by the street network, and gives rise to heterogeneous patterns of activity.

The extent to which the behavioural principles described above can be applied to policing remains an open question. Police officers represent a singular case, since a significant proportion of their activity is undirected (as during generic patrol, when not responding to an incident). Furthermore, their activity is not constrained by traffic laws or road conditions in the same way as it is for civilians, since these typically do not apply under emergency (*i.e.* blue-light) conditions. There is reason to expect, therefore, that officer behaviour will differ from that of the general public, while also being influenced by the spatial distribution of demand. The difficulty of obtaining police movement data, however, means that these issues remain largely unexplored, with only a few empirical examples in the literature (*e.g.* ??). Furthermore, even if such routing behaviour was well-understood, the question of how this translates to macro-level coverage would still present a significant challenge. It is this issue that we seek to address directly in this work.

1.4. *The role of ‘anchor points’*

Nodes or specific locations to which routine behaviour are tied are a central concept of crime pattern theory (?). The movement of people between these nodes as part of their routine activities is the context in which opportunities for crime present themselves (?). There is a large literature on the relationship between different types of nodes and the generation of crime. For example, many types of facility have been linked to concentrations of crime problems and have therefore been labelled as ‘crime generators’, ‘crime attractors’ or ‘risky facilities’ (see ?).

The literature on significant places in the commission of crime offences has also extended to journey-to-crime research. Here, it is proposed that offenders, all else being equal, will have a preference for shorter journeys and therefore will be more likely to offend nearer to their home locations. This is linked to the concept of familiarity decay

(?). In a nutshell, the least effort principle, along with a preference for known locations, should lead to a dramatic drop in the number of crimes committed at increased distance from home. There is an extensive body of empirical research confirming this (see ?). Extending these principles to policing - another occupation in which travel is necessary but routing and end point are (sometimes) discretionary - it seems logical that similar principles could be applied. Hence, the location of police stations should have a significant influence over the distribution of travel patterns as this will be the point of origin of those movements. To our knowledge this has not been addressed by previous research, but we suggest that the degree to which demand is met by vehicular policing patrol will, in part, be a function of the travel distance involved.

In summary, we are interested in: (1) investigating the relationship between supply and demand in policing; (2) assessing whether this is mediated by the form of the street layout; and (3) assessing whether the location of police stations influence the balance of provision.

2. Methods

As introduced in the previous sections, the aim of our work is to examine the alignment between the spatial distribution of police presence and that of the demand for police service. In the following two sections, we will present our empirical analysis, which is based on data for the city of London, UK. We will describe the data used, along with the pre-processing steps required to facilitate the primary analysis. We will then show the disparity between supply and demand for a number of London boroughs, through which a number of distinctive trends can be identified. Following this, we will use a quantitative model to examine how our measures of police coverage relate to a number of features of the urban environment, including the structure of the street network and the location of police stations.

2.1. *Data and pre-processing*

In order to perform our analysis, the two key phenomena of interest - police presence and calls for service - must be measured in such a way that they can meaningfully be compared. Our analysis is therefore based on two datasets for the city of London which relate to the same interval of time: namely, the 5 month period between the 1st of June and 31st of October 2015. These datasets were provided by the Metropolitan Police Service, and provide a complete record of activity in the 32 boroughs of London. The time window examined was determined by data availability, as these were the only months for which concurrent data were available; while it is possible that seasonal effects may be present, we have no reason to expect that these months are not representative of the wider picture, in the terms that we measure. Before describing the movement and call datasets in detail, we will first introduce the street network data which provides the spatial setting for our analysis.

2.1.1. *Street network*

The network data used throughout this work is provided by Ordnance Survey (OS), via the Integrated Transport Network layer of its MasterMap product. This dataset contains records for all streets within the boundary of London which are suitable for vehicular travel. In each case, the geometry of the street it provided, along with

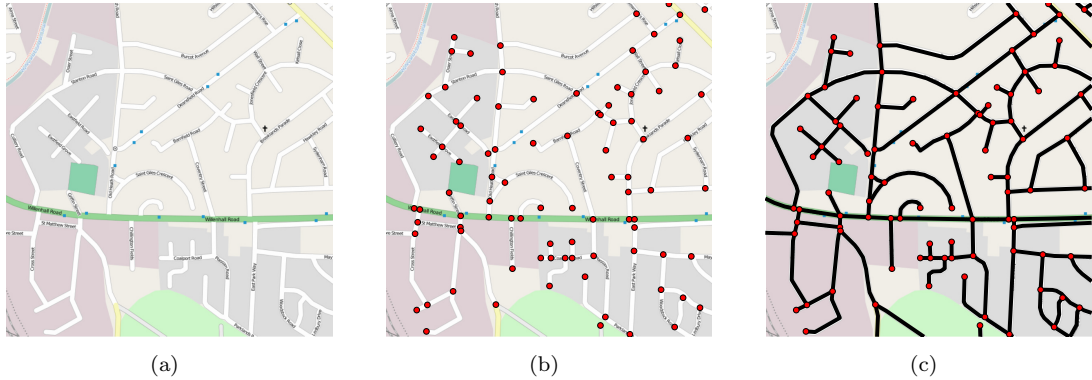


Figure 1. The construction of the primal representation of a street network: a) the original map, b) nodes placed at each junction, c) links added between any pair of junctions connected by a street segment.

contextual information such as road type, administrative classification and routing restrictions (e.g. whether a street is one-way).

In order to measure the structural properties of the network, as required for the latter part of our analysis, it is necessary to represent the data in formal mathematical terms. This can be done using terminology from graph theory (see ?), which concerns the structure of pairwise relationships between discrete entities. Formally, a *network* $G = (V, E)$ is a set of *vertices*, V , together with a set of *edges*, E , which connect pairs of vertices. In basic terms, a network is simply a collection of objects, some pairs of which are linked: vertices that are connected by an edge are said to be *adjacent*. The degree of a vertex is defined as the number of other vertices to which it is adjacent.

The most natural way to express a street network in this form is to take vertices to represent junctions (*i.e.* points at which roads intersect) and to place an edge between any pair of junctions that are directly connected by a street. Edges therefore represent street segments: the sections of road between immediately-neighbouring junctions¹. This is referred to as the ‘primal’ representation (?), and an example of its construction is shown in Figure ???. Although other representations exist (??), the primal approach is taken here because it is the only one that preserves geo-spatial information: metric distance, for example, is not well-defined in other representations.

2.1.2. Demand for service

The demand for policing can be defined in a number of ways, and the choice of which to use is dependent upon context. In the majority of criminological research, for example, demand is defined from the perspective of crime prevention, and is therefore measured in terms of volume of criminal incidents. In this work, however, our focus is broader: we are concerned with policing as a whole, and the nature of overall service delivery. From an operational perspective, the demand for policing includes not only crime-related activity but also the various other functions which officers are required to perform. In practice, all situations which require the physical presence of officers constitute demand, and the full range of police activity must therefore be considered in order to properly evaluate resource allocation.

With the above reasoning in mind, the analysis we present here is based on call for

¹The raw data contains elevation information, which means that cases where two streets cross each other but no junction is present (e.g. bridges) can be correctly represented. In such a case, no junction is present between the crossing streets.

service (CFS) data. In particular, we use records from the Computer Aided Dispatch (CAD) system operated by MPS, which is the primary tool used to manage officer deployments in real time. The data we use covers the period between the 1st of June and 31st of October 2015, and contains a record for every incident to which officers were deployed within London during this period². The incidents in question can be generated in response to emergency calls or by other means, and include events of many types: crimes, public order situations and responses related to public health are all represented. We do not distinguish between call types; that is, all calls are treated equally and are not weighted according to, for example, seriousness. Our rationale for this relates to our conceptualisation of demand as the need for police response: although some incidents will undoubtedly require more police time once on the scene, at the point of dispatch all calls require a vehicular response, and therefore represent an equivalent draw on the ‘supply’ of policing. In addition to this, the variation in police time across incidents is likely to be so great that any attempt to control for this would present significant challenges of its own.

The location at which each event was generated is provided in terms of British National Grid coordinates. In order to facilitate our network-based analysis, we associate each incident with the closest street segment to which it took place. In order that we can be confident that each of these incidents did indeed take place at a location on the network (rather than in a park, for example), we impose a maximum distance of 50 metres on this association; incidents located further than this from the street network are discarded from the analysis. After this processing step, which typically results in 1-2% of incidents being discarded for each borough, we have a dataset consisting of 1,194,255 incidents occurring across London, each of which is associated with a particular street segment.

2.1.3. Police movement

The measure of police presence used in this study is based on the analysis of vehicle movement data provided by MPS. The focus on vehicular travel is partly motivated by practical considerations (most prominently data quality) but also reflects the dominant mode of policing in London. Under normal conditions, the majority of response work is carried out by vehicle-based officers, and it is those units to which dispatches will be assigned in the first instance. In terms of response, therefore - and, to an extent, the provision of visible police presence - the positioning and routing of vehicles is the central issue for resource allocation.

In London, all police vehicles are equipped with Global Positioning System (GPS) sensors, and it is this feature which allows their movements to be examined at an unprecedented level of detail. Each vehicle is equipped with a mobile data terminal, which records the position of the vehicle, while active, approximately every 15 seconds. Each entry contains a combination of positional data (coordinates, bearing and speed) and tasking information (status, callsign and a reference to any incident to which it is currently assigned), which can be used to reconstruct the historical activity of each unit. In our analysis, we examine the activity across our 5 month study period of all vehicles operated by MPS.

In terms of positional information, these GPS logs require significant pre-processing in order to be suitable for analysis. Since the purpose of the study is to carry out analysis at the micro level, it is crucial that the locations of police vehicles are known to high levels of accuracy and completeness. Two aspects of the movement dataset, how-

²Duplicate records, representing multiple calls or dispatches to the same incident, have been removed.

ever, have the potential to compromise this: the error associated with GPS readings, and the issue of sampling.

The first of these is due to the fact that GPS readings are subject to measurement error, which means that the recorded locations may be inaccurate, and indeed may not lie on the street network. The most immediate solution to this, by which each point is associated with the street segment to which it is closest, leads to frequent mis-assignments (particularly in dense street networks such as those found in London) and is therefore inadequate.

A further problem concerns the fact that the recorded GPS locations provide only a sample of a vehicle’s positions. If only recorded locations were used, this would fail to record visits to intermediate segments; *i.e.* segments traversed on the journey from one recorded location to the next. On the other hand, situations in which a vehicle was stationary for an extended period would contribute multiple readings to the same segment, whereas in fact only one visit had taken place.

In order to overcome these issues, it is necessary to convert each vehicle’s GPS log to an estimate of the full path which it took. This is a process known as ‘map-matching’ and is most commonly performed in applications such as route planning or satellite navigation systems. Though no method can be considered definitive, a diverse range of approaches have been published (?), including those based on the minimisation of Fréchet distance (??) or using Kalman filters (?).

In this study, we employed a customised implementation of the algorithm introduced by ?. The first step in this algorithm is to establish the possible locations on the street network at which each GPS record could have been generated. For each GPS record in the sequence, the m closest points on the street network are found, with a weighting assigned to each according to its distance from the recorded location. Given these sets of candidate points for each record, the next step is to evaluate the possible journeys between them. For each pair of consecutive records, the shortest paths (taking into account routing restrictions) between all pairs of candidate points are found. Each of these paths is assigned a weighting according to its length, relative to the Euclidean distance between the observed GPS locations. The final step is then to select the candidate path with highest weight (taking into account both observation and transition weights): this is equivalent to selecting one candidate point for each observed location. The output of this algorithm is a full path through the network.

It is worth emphasising here that the vehicle presence is a single measure and it was not possible to distinguish between different types of trip or operational function. Hence, we do not for example compare those cars dispatched for calls for service with those on proactive patrol. Neither do we count any of the time that a car is stationary on a street segment. Whilst readers might see this as a limitation, it is an appropriate measurement strategy in this case. This is because we are not testing the efficacy of policing activity, we are purely discerning the degree to which police vehicles are visibly active in areas with a high demand for services. So, the measure here can perhaps be best understood as a representation of supply from a public perception viewpoint.

2.2. Comparing supply and demand

The datasets described above, once processed, provide us with complete information concerning both the supply of policing (in the form of vehicle presence) and the demand for it (in the form of calls for service). The aim of our work is to explore the alignment

between these quantities. In order to do this, it is necessary to first establish at what scale, and in what terms, these factors can most meaningfully be compared.

From an organisational perspective, the MPS is divided into 32 Borough Operational Command Units, corresponding to the 32 boroughs of London. These units have their own dedicated officers, and are largely autonomous in terms of their day-to-day operational practices. In particular, it is at the borough level that the majority of fine-level resource allocation decisions are made: the prioritisation of particular areas and offences, along with general procedural direction, is determined by borough sergeants. It is at this level, therefore, that the distribution of activity can be most meaningfully understood. Furthermore, any analysis aggregated over multiple boroughs would fail to account for differences in the level of resource available and the nature of demand. In our analysis, we therefore consider all boroughs independently.

After partitioning both datasets by borough, the next stage is to compare their values on each street segment. Since the units are different in each case - ‘total minutes visited’ in the case of supply, and ‘total number of crimes’ in the case of demand - they must first be converted to a form which can be directly compared. We do this by normalising the values in each case according to their borough level totals. In particular, for a given edge e , we define the *supply*, denoted s_e , as

$$s_e = \frac{\text{seconds of vehicle presence on } e}{\text{total seconds of vehicle activity across borough}}. \quad (1)$$

Analogously for the number of calls for service, we define the *demand* on e , denoted d_e , as

$$d_e = \frac{\text{number of calls arising on } e}{\text{total calls across borough}}. \quad (2)$$

In simple terms, these can be interpreted as ‘the fraction of police time devoted to e ’ and ‘the fraction of calls generated by e ’ respectively. Our interest is in the extent to which s_e and d_e are unequal for each borough (equality being what would be expected under a naïve assumption of perfect allocation).

Since s_e and d_e are both non-dimensional, they can be compared directly. To capture the disparity between them, we therefore define the quantity *provision* as the difference between the two, so that the provision, p_e , for each edge, e , is

$$p_e = s_e - d_e. \quad (3)$$

Provision therefore captures the extent to which the police activity devoted to a street segment is higher than is warranted by the proportion of calls generated. Positive provision therefore corresponds to a situation in which a segment receives more police presence than it would on the basis of calls alone, and *vice versa*. This measure therefore quantifies the disparity between the two key phenomena, and will be the main focus of our analysis.

2.3. Analytical approach

The quantity ‘provision’ is the basis for our quantitative analysis of policing activity. Since it provides a measure of the alignment between police resourcing and demand

for every street segment in greater London, it can be analysed in several ways, and we present a number of perspectives.

Our initial analysis is primarily descriptive, examining the properties of the ‘provision’ variable itself and its spatial distribution. Motivated by the features of these maps, we then seek to build a statistical model of provision, relating it to amongst other things the structural features of the street network. To this end, we introduce the network metric betweenness, which acts as an estimate of the usage of street segments, and estimate ordinary least squares models of provision. Finally, in order to examine the operationally-significant issues of under- or over-provision, we examine the residuals of these models and their spatial distribution.

As noted above, our analysis is performed on a borough-by-borough basis, and accordingly we will present numerical results for all 32 boroughs. In addition, in order to illustrate particular points - by showing maps, for example - we will do this by presenting results for selected individual boroughs.

3. Results

In this section, our examination of the provision variable is presented, building from descriptive analysis to statistical modelling.

3.1. *Distribution of provision*

Before examining its spatial distribution, it is instructive to examine the overall distribution of provision. By doing this, we can see the typical range of values that it takes, and check whether any extreme values can be observed.

Since it is defined as the difference between two normalised quantities, provision is guaranteed to have a mean of zero across each borough. Our interest is therefore in its dispersion around zero, and any degree of skewness that can be observed.




Figure 2. Histogram showing the distribution of provision, p_e , across all street segments in the London borough of Brent. The distribution displays a clear peak around 0, and is approximately symmetric around this value.

Figure ?? shows the distribution of provision for the borough of Brent. A number of general features are immediately apparent. On one hand, it is clear that the distribution of provision is broadly symmetric about 0, suggesting that positive and negative provision are balanced in terms of their overall character. In addition to this, it is clear to see that the distribution is strongly peaked at zero, with a majority of values lying in the range $[-0.0001, 0.0001]$. This implies that, on the whole, the difference between supply and demand measures is very small for most segments in the borough. Nevertheless, it is also apparent that some segments give rise to values which are much more extreme: a number of cases with absolute values greater than 0.0005 can be seen, for example. These cases are likely to be interesting from a theoretical and practical point of view, since they imply a substantial mis-alignment between supply and demand.

3.2. *Spatial distribution*

While the overall distribution of provision is informative in a general sense, our primary interest concerns its distribution in space. It is by examining this that we can reveal whether the extreme values observed above occur in any discernible pattern, and whether provision exhibits any spatial biases. As well as offering insight into the behaviour of police officers and their resource allocation practices, this also has the potential to be of practical value by highlighting areas of (potentially adverse) imbalance.

In Figure ??, we show maps for two individual boroughs - Brent and Merton - in which street segments are coloured according to provision. For visual clarity, we show two maps in each case: one showing positive values of provision (in which negative values are shown as zero), and one showing negative values (with positive values set to zero). On the left, therefore, bright segments are those on which supply is greater than demand, while the opposite is true on the right.

Looking first at positive provision, it is apparent that high values of provision appear to occur along extended ‘main’ routes through the boroughs. Although the types of roads are not shown, the routes appear to form an arterial system such as would be formed by the primary roads in each borough. That they should show high positive provision is therefore not unexpected: in a hierarchical road system, it is typical for traffic to be ‘funnelled’ towards major roads and for these to therefore be used more heavily. The fact that provision is high on these segments (and therefore that supply is high) implies that the movement of police officers is consistent with this: their travel is heavily focussed on a small number of primary roads.

In the case of negative provision, a much less regular pattern is observed: bright segments are scattered throughout the boroughs, and do not appear to form any clear extended routes. In order to interpret the patterns, it is worthwhile to reflect on what negative provision implies: that demand is greater than supply. Bright segments in these maps must therefore have generated a non-zero number of calls, and received a volume of police presence which was not commensurate with that. Their spatial distribution can largely be rationalised as a consequence of the positive provision on major streets: because they spend more time on those, they must spend less on segments which lie away from primary routes. When these segments generate demand, therefore, this is manifested as an imbalance.

Put in simple terms, the results collectively suggest that the presence of police (*i.e.* supply) is biased towards central primary routes to a greater extent than is warranted by the distribution of demand. While this is true in basic terms, however, this does not

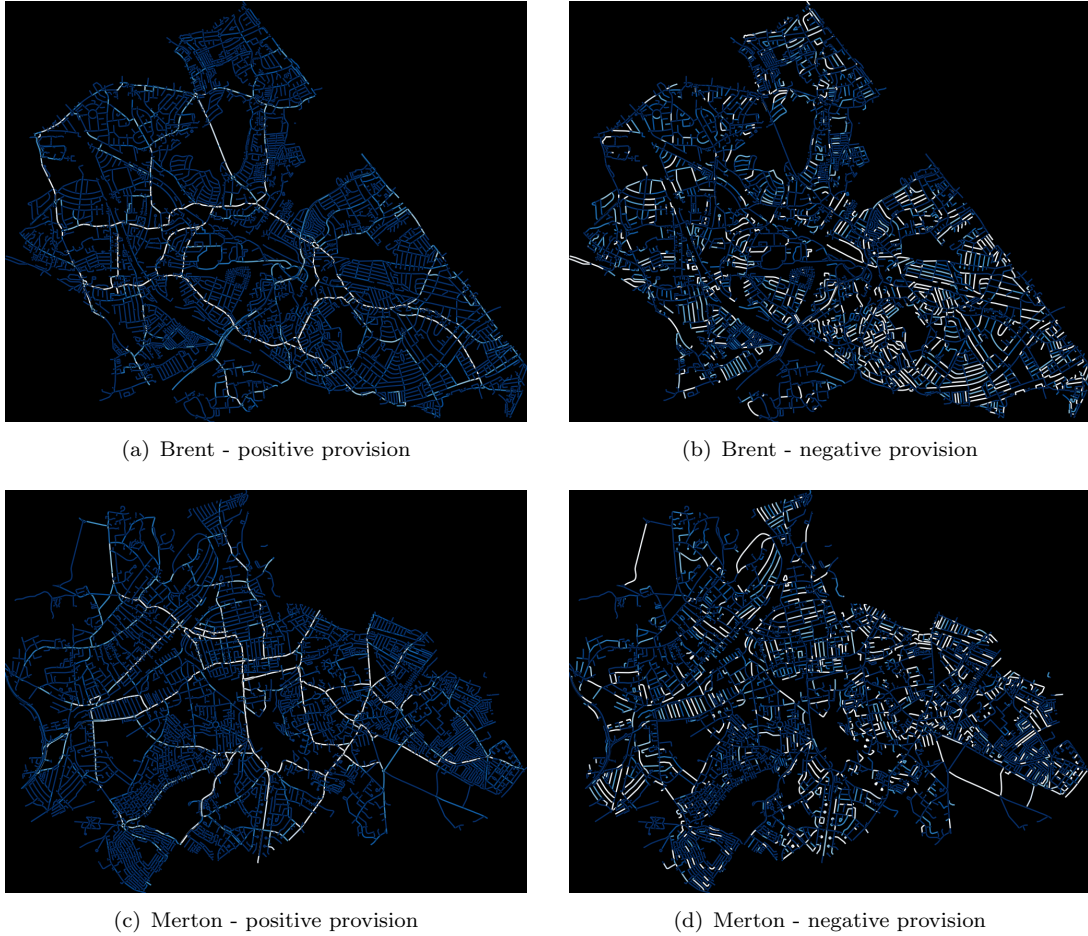


Figure 3. Positive and negative provision in Brent and Merton.

imply that resources are allocated inefficiently or that officers behave sub-optimally. There are valid reasons why primary routes should be favoured: their purpose is to act as shortcuts between areas, and it is rational that officers seeking to complete journeys as fast as possible should favour such routes. In order to perform more nuanced analysis, therefore, it is necessary to account for aspects of urban structure.

3.3. *Statistical model of provision*

It is clear from the maps presented in the previous section that a relationship exists between provision and network structure. At this stage, this can only be expressed in qualitative terms - ‘more major streets display higher provision’ - and the sense in which the centrality of streets is defined is somewhat vague. Network analysis, however, provides a means by which concepts such as these can be formalised and quantified in a systematic way. Using this approach, in this section we seek to model the relationship between network structure and provision.

As described in Section ??, the centrality of streets can be defined in a number of ways, corresponding to different aspects of their role in the network. One of these which is particularly relevant in the context of travel through the network is *betweenness*, which measures the extent to which a street segment acts as a conduit for movement

on the network. More specifically, betweenness measures the frequency with which street segments feature in paths through the network, and therefore represents a proxy estimate of the volume of traffic that they will carry. It is calculated via the following process, in which B_e is taken to denote the betweenness of each edge e :

- (1) for all edges, initialise $B_e = 0$
- (2) consider all pairs of vertices, v and w
- (3) in each case, find the shortest path(s) through the network between v and w
- (4) every time an edge appears on one of these paths, increment its B_e by $\frac{1}{m}$, where m is the number of shortest paths between v and w (m will only be greater than 1 in rare cases where multiple distinct paths have exactly the same length).

This can be expressed formally by introducing some further notation. If σ_{vw} is defined as the total number of shortest paths between two vertices v and w , and $\sigma_{vw}(e)$ is defined as the number of these paths which pass through e , then the betweenness of an edge e is given by

$$B_e = \sum_{\substack{v,w \in V, \\ v \sim w}} \frac{\sigma_{vw}(e)}{\sigma_{vw}}, \quad (4)$$

in which \sim denotes ‘there exists a path between v and w ’ (this simply restricts the sum to pairs for which $\sigma_{vw} = 0$).

One further refinement can be made by imposing a limit on the length of paths considered in the betweenness calculation. To do this, a given maximum radius, r , is specified, and only paths of length less than r contribute to the betweenness calculation (*i.e.* step (4) above is only performed if the paths found in (3) are shorter than r). Using such a modification has two primary benefits. On one hand, it removes the possibility of edge effects, where misleading values are produced near the extremities of the study area. In addition, by removing the contribution of long-distance trips, it produces a more ‘local’ measure of betweenness, which reflects more closely the patterns generated by typical travellers. In the remainder of the text, we use betweenness values based on a radius r of 1,000m (other values were tested and gave qualitatively similar results).

Figure ?? shows the distribution of betweenness in Brent, the same borough as was shown in Figures ?? and ?. It is clear from these maps that betweenness reflects an intuitive sense of centrality: high-valued segments form an artery-like system, and the brightness decreases through the hierarchy. In addition, visual comparison of Figures ?? and ? reveals striking similarities: the majority of segments on which high positive provision is observed are also of high betweenness. Motivated by this similarity, we attempt to quantify this relationship using a regression model.

In line with the overall structure of our analysis, we build individual ordinary least-squares regression models for each borough in London, with our measure of provision as the dependent variable. In addition to betweenness, we also include two other street segment attributes as independent variables. The first of these is road type: the categorical classification assigned by OS (*e.g.* ‘A Road’, ‘B Road’), which corresponds broadly to the intended use of the segment. This is dummy-coded as a categorical variable, using ‘A Road’ (the highest class) as the reference value.

We also include a further variable which measures, for each street segment, the dis-



Figure 4. Maps of the street network of Brent, coloured according to: a) betweenness, and b) station proximity.

tance to the closest police station. The reason for including this is that segments which are closer to stations will, in general, receive greater coverage purely as a result of their proximity to a ‘base’: all vehicles must begin and end their journeys at a station, and may return to them during a shift if operationally required. This variable is therefore included to control for this effect, and account for the bias towards neighbouring streets. Figure ?? illustrates this proximity measure for the borough of Brent, which contains two stations.

The fitted coefficients of this regression model are given in Table ?. Considering first our primary issue of interest, these results provide strong support for the hypothesis that provision is positively associated with betweenness. The coefficient in question is positive for all boroughs, and significant at $p < 0.05$ in all cases other than Westminster³. All boroughs give rise to broadly similar values, although those from Kingston-upon-Thames and Barking & Dagenham are notably high; it is possible that the particular topologies of these boroughs give rise to a particularly pronounced effect. In general, though, the results are as anticipated: provision is higher on segments with higher betweenness. This is likely to be a result of such segments being favoured in route-finding and therefore receiving excess supply.

The results for the other control variables are also broadly in line with expectation. The influence of station proximity is negative and significant in most cases, which accords with the hypothesis that segments which are closer to stations will naturally receive greater provision. There are, however, a number of cases where the effect is non-significant: this may arise from the particular location of stations within the boroughs in question, or the distribution of officers between them. The categorical variables corresponding to road classification are significant in all cases, suggesting that these have an influence over and above that of betweenness. Variations in speed limits and road capacity provide a plausible explanation for this, since these factors also affect decision-making and route planning. In terms of the categories themselves, the results generally follow the hierarchy of the classification: provision is highest on A Roads, followed by B Roads, and so on. The only exception to this is the inverted position of ‘Local Street’ and ‘Alley’: in all cases, alleys have higher provision. One explanation

³Although the issue of multiple comparisons is relevant when estimating an array of models such as this, these results remain significant after applying a Bonferroni correction in all cases other than that of Camden

Table 1. Coefficient values for the ordinary least squares model of provision described in the text, fitted independently for each borough in London. The superscript * indicates that $p < 0.05$ for the corresponding coefficient.

Borough	$\log(B_e)$	Station proximity (km)	B Road	Minor Road	Local Street	Alley
Barking And Dagenham	0.06*	-0.01	-0.40*	-0.48*	-0.83*	-0.71*
Barnet	0.01*	-0.01*	-0.11*	-0.08*	-0.23*	-0.19*
Bexley	0.02*	-0.00	-0.15*	-0.22*	-0.40*	-0.30*
Brent	0.01*	0.00	-0.20*	-0.28*	-0.40*	-0.32*
Bromley	0.01*	-0.00	0.03	-0.26*	-0.35*	-0.30*
Camden	0.01*	-0.02*	-0.01	-0.14*	-0.38*	-0.30*
Croydon	0.01*	0.00	-0.25*	-0.34*	-0.40*	-0.34*
Ealing	0.02*	-0.01*	-0.11*	-0.28*	-0.40*	-0.33*
Enfield	0.02*	-0.01*	0.02	-0.17*	-0.33*	-0.24*
Greenwich	0.02*	-0.01*	-0.20*	-0.20*	-0.40*	-0.30*
Hackney	0.02*	-0.05*	-0.10*	-0.16*	-0.46*	-0.32*
Hammersmith And Fulham	0.03*	-0.04*	-0.01	-0.33*	-0.63*	-0.44*
Haringey	0.02*	-0.03*	-0.06	-0.26*	-0.53*	-0.36*
Harrow	0.02*	-0.00	-0.28*	-0.22*	-0.45*	-0.34*
Havering	0.03*	-0.01*	-0.33*	-0.25*	-0.52*	-0.41*
Hillingdon	0.01*	-0.00	0.08*	-0.25*	-0.34*	-0.30*
Hounslow	0.01*	-0.00	-0.22*	-0.27*	-0.38*	-0.29*
Islington	0.01*	-0.07*	0.35*	-0.20*	-0.38*	-0.23*
Kensington And Chelsea	0.03*	-0.08*	0.12*	-0.22*	-0.48*	-0.38*
Kingston Upon Thames	0.06*	0.00	-0.49*	-0.53*	-0.71*	-0.54*
Lewisham	0.02*	-0.02*	0.02	-0.23*	-0.37*	-0.26*
Merton	0.03*	-0.01*	-0.29*	-0.36*	-0.52*	-0.41*
Newham	0.02*	0.00	0.04	-0.27*	-0.44*	-0.34*
Redbridge	0.03*	-0.01*	-0.11*	-0.09*	-0.37*	-0.26*
Richmond Upon Thames	0.03*	-0.02*	-0.38*	-0.49*	-0.58*	-0.45*
Southwark	0.01*	-0.01*	-0.07*	-0.07*	-0.30*	-0.26*
Sutton	0.02*	-0.02*	-0.12*	-0.52*	-0.69*	-0.59*
Tower Hamlets	0.02*	-0.05*	-0.14*	-0.26*	-0.43*	-0.37*
Wandsworth	0.02*	-0.02*	-0.16*	-0.29*	-0.48*	-0.37*
Westminster	0.01	-0.02	-0.11*	-0.13*	-0.23*	-0.21*

for this is that few calls for service are geo-coded to alleys, since they typically have few associated landmarks or address points. This would cause the demand on these segments to be artificially low, therefore increasing their provision. Nevertheless, it cannot be discounted that alleys do indeed receive more coverage in the course of standard officer behaviour.

3.4. *Residuals of provision*


The results of the previous section are informative, in the sense that they confirm that mis-alignment between supply and demand is, to some extent, a function of network structure. This is a positive result, and confirms that urban structure plays a role in shaping coverage. The implications of this from a practical perspective are, however, limited, since the model simply encodes the (relatively unsurprising) notion that officers tend to favour more central streets when navigating. For street segment provision levels to be consistent with this therefore indicates only that they can be explained by this mechanism. In fact, therefore, of greater significance is the question of where this consistency is *not* observed; that is, in which locations our results deviate from the model. With this in mind, in this section we examine the residual values of the fitted model.

The residual for each street segment quantifies the extent to which it deviates from that predicted by the fitted model: a positive residual indicates that provision is greater than would be expected, and *vice versa*. The interpretation of this depends, to some extent, on the predicted value, since this indicates whether under- or over-provision would be expected. If, for example, both predicted value and residual are positive, this implies that provision would be anticipated to be positive, but not to this extent; this is over-provision beyond what can be explained by topological factors. Conversely, for both to be negative indicates that a segment receives even less provision than its non-centrality would imply. With this in mind, we examine four cases, corresponding to positive and negative values of both predicted values and residuals.


In Figure ??, we show these cases for the borough of Brent (other boroughs show similar results, in the sense than the same qualitative observations can be made). Although Figures ?? and ?? reveal relatively little, striking results can be seen in Figures ?? and ??, which correspond to the cases described in the previous paragraph.

Figure ?? shows a number of high negative residual values scattered throughout the borough. Since these are cases in which provision was expected to be negative (*i.e.* supply is low-valued), it is likely that these correspond to sites of particularly high demand. In isolated cases, therefore, the presence of these is unsurprising and likely to be caused by a burst of incidents. Where such segments cluster, however, may indicate a cause for concern: such areas generate high volumes of incidents, and are consistently under-supplied in terms of police presence. From the perspective of police resource allocation, this suggests that this may be caused by an underlying systematic factor, and may warrant remedial action. An example of such an area can be seen in the central south area of Figure ?. It is worth re-emphasising here that the residual under-provision referred to here is additional to that which can be accounted for by differences in network structure and police stations proximity.

At the other extreme, Figure ?? highlights a number of segments which receive even greater provision than can be accounted for by our structural model. This suggests that there is an additional bias towards these segments, which appear to be a defined



(a) Negative predicted, negative residual



(b) Negative predicted, positive residual



(c) Positive predicted, negative residual



(d) Positive predicted, positive residual

Figure 5. The street network of Brent, with segments coloured according to residuals relative to the fitted model for provision.

subset of the more central ones, causing them to receive disproportionate volumes of police attention. This could be a result of habit on the part of officers, or could reflect other factors which are not captured by our measure of provision. In either case, this represents a further opportunity for behaviour change. The additional provision on these segments could be considered to be ‘spare’ capacity - not warranted by demand or topological factors - which might be more effectively deployed elsewhere. One possible course of action would be to encourage officers to take alternative routes when travelling in these directions, using less central segments when possible.

4. Discussion

The optimal spatial distribution of police resourcing has been the subject of both academic and political interest. Past evidence from empirical analysis has been limited to either studies at the aggregate level - for example assessing the impact of police force size on crime reduction (?, offer a recent review) - or investigating the relationship between police activity and crime reduction as a result of special intervention, often in focused geographical areas (?). The research presented in this paper was motivated by

the need to understand a distribution distinct from either of these: the everyday provision of policing at the micro-scale. Such research is only possible with the availability of spatio-temporal tracking data collected through GPS devices and with sufficient technical understanding of the resulting data to reliably prepare it for analysis.

In the interest of addressing community needs concerning security services, we chose to focus on the demand for police service more generally rather than focusing on the relationship between police presence and levels of crime *per se*. The principles of service delivery based on community priorities and reassurance agendas are firmly established and it is, quite simply, important that the police go where the public ask for them, within the limits to which this is plausible. Furthermore, it is crucial that this is done visibly, where possible. Previous research has confirmed that higher police visibility can result in improved perceptions of the legitimacy of the police (??) and in reduced levels of fear (?), and so desired outcomes should not be restricted to crime reduction alone.

The results of our research, which focus on police vehicle patrol in the area of Greater London served by the Metropolitan Police Service, reveal a number of key findings. First, positive and negative provision, where demand and supply do not align, exists, and this is not spatially random. Police vehicles have a preference for main, better-connected, roads and will choose these routes over and above other alternatives that have the same level of request for police time. This is hardly surprising given that police driving through streets have a number of purposes including the use of connected streets to get them quickly to other places. Second, provision which is over and above demand is predicted by a number of characteristics of the street network; segments with high betweenness and hence greater usage; those with more major road classifications, and those which are in closer proximity to police stations all have a tendency to be over-patrolled by vehicles. Third, and possibly of most interest to operational policing, there is still over- and under-provision on some segments when network characteristics are controlled for. In these cases preference or avoidance cannot be attributed directly to accessibility.

This last results suggests a behavioural bias in collective police vehicle patrolling that could be useful to address in making the distribution of resourcing as ‘optimal’ or ‘efficient’ as possible. From a visibility point of view, addressing this bias is perhaps most significant in areas where the ‘residual’ over- or under-provision clusters in space. It seems likely, for example, that a cluster of streets with unusual under-provision will have some features that make them less attractive to police patrol. These could be associated with land use: perhaps the streets have a large number of risky facilities (?), making them more uncomfortable places to visit. On the other hand, they might lack the facilities that make vehicular policing more comfortable as a day-to-day activity. Another explanation relates to the reputation of areas. These areas could, for instance, be those which are generally regarded as ‘no-go’ zones by non-residents for anything other than essential service provision. Again, the alternative is also viable; that they are areas regarded as less of a problem, even given a certain volume of calls for service. These might, for instance, be places with higher levels of collective efficacy (?) where the public are more able to self-police.

To assess these candidate hypotheses more detailed research would be required, quite possibly of an ethnographic nature in addition to an analytical one. Whilst outside the scope of this initial exploration, we would recommend future research to unpack the cause of both under- and over-provision that is not attributable to accessibility. We also acknowledge that there are some elements of the analysis where assumptions have been necessary. First, we have not distinguished between types of police vehicle,

the speed they were driving at or whether they were on ‘blue lights’ or not. Second, we treat each call to the police equally and do not weight by type of incident. As mentioned above, this is a conscious choice: our interest is in the trips and routes that police vehicles take, and their consequent availability to react to demand. For this aim, distinguishing between type of incident and type of response is not necessary: at the point of dispatch, all incidents will require a response, independently of the form and duration that the demand ultimately takes. Nevertheless, we acknowledge that there will be exceptions in the data: some calls might not elicit a vehicular response, for example, while other situations might occur where multiple calls are elicited for a single incident. This is particularly likely to happen where an event is more serious, such as a violent crime or a large-scale public order emergency. In these cases a one-to-one assumption of one vehicular deployment for one call might not stand.

In addition, police vehicles have many functions and here we have focused on their ‘visible moving presence’ function alone as a way of conceptualising supply. This means that we do not account for the presence of stationary cars in certain locations. Stationary police cars might have the potential to demonstrate available resource, although that we suspect that parked non-manned vehicles are unlikely to turn as many heads as moving ones. Neither do we consider policing resource external to vehicles, either in the form of foot patrol or any duties relating to response to incidents. We would argue here, however, that to be useful in terms of operational police resourcing, the advice offered should be function-specific. This makes it far more manageable to effectively intervene. By means of a parallel, evaluation research has advocated the assessment of particular prevention strategies with particular goals, mechanisms and contexts and would not advise more general approaches (?).

A final limitation is a challenge faced by any researcher examining the impact of policing activity on crime. This is the problem of potential endogeneity, where the causal direction of the relationship between variables cannot be assumed (see ?). In this case, it is possible that the presence of police itself results in lower demand (since the supply is already present), or indeed vice versa. This is certainly true, and feedback effects such as these limit what can be concluded from this work about the underlying relationship between supply and demand. However, the focus of this study is on the alignment of policing activity with demand as it is observed: from the perspective of the police, even though call-for-service data is subject to these biases, it is the basis on which they must judge demand (and the terms in which their performance is judged). In this work, we are not attempting to *explain* the direct relationship between the response and levels of calls, but instead to identify any manifesting imbalance between them. This will, however, be an avenue for further research.

Independent of the political debates concerning police resourcing, it is a truth that to adapt to the current austerity required the police have had to focus on methods that increase their efficiency. Little research has looked at how police officers actually spend their time in practice or on the efficiency of the police operational resourcing practices. As one of the initial steps in understanding how demand from the public is serviced, in this article we have proposed a method for understanding the geography of police provision. Specifically, we examine the geography of over and under- provision. That is, parts of the street network that have an over or under representation of vehicular police ‘presence’ on the basis of demand. Maps such as Figure ?? could be useful operationally: for example, if produced regularly and disseminated, police drivers could, where possible, plug gaps in provision. This is likely to be easier to do on non-time-dependent routine visits or proactive patrol, which itself could provide structure and focus to officer time: as identified by ?, a lack of direction to discretionary time can

lead to substantial inefficiencies within policing. We suggest that a movement towards automated systems that can assist with route decision making – such as a route guidance system which allowed directions via routes optimised for ‘provision balancing’ – could even rebalance spread without much additional resource. Re-balancing provision, as well as ensuring vehicles service places of need, has the potential to lead to an increase in satisfaction with the police as perceived by the public. If the police are seen as being available in places with current demand, they will be viewed as responsive and proactive. It is therefore useful for decision makers to understand collective bias in patrolling. Of course, there needs to be appropriate sensitivity to individual privacy, freedom and variation. The kind of analysis undertaken here only makes sense at the aggregate level and we advocate an approach that would call for collective action in areas of concern with suggested routes for when this is viable, rather than prescribed ones.

References

- Alt, H., *et al.*, 2003. Matching Planar Maps. *In: Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '03, Philadelphia, PA, USA. Society for Industrial and Applied Mathematics, 589–598.
- Ariel, B., Weinborn, C., and Sherman, L.W., 2016. “Soft” policing at hot spots—do police community support officers work? A randomized controlled trial. *Journal of experimental criminology*, 1–41.
- Armitage, R., 2007. Sustainability versus Safety: Confusion, Conflict and Contradiction in Designing Out Crime. *In: G. Farrell, K.J. Bowers, S.D. Johnson and M. Townsley, eds. Imagination for Crime Prevention: Essays in Honour of Ken Pease*. Crime Prevention Studies, vol. 21. Monsey, NY: Criminal Justice Press, 81–110.
- Audit Commission, 1996a. *Streetwise: Effective Police Patrol*. London, UK: Audit Commission.
- Audit Commission, 1996b. *Tackling Patrol Effectively*. London, UK: Audit Commission.
- Beavon, D.J.K., Brantingham, P.L., and Brantingham, P.J., 1994. The Influence of Street Networks on the Patterning of Property Offences. *In: R.V. Clarke, ed. Crime Prevention Studies*. vol. 2. Monsey, NY: Criminal Justice Press, 115–148.
- Bevis, C. and Nutter, J., 1977. *Changing street layouts to reduce residential burglary*. Minneapolis, MN: Minnesota Crime Prevention Center.
- Bollobás, B., 2002. *Modern Graph Theory*. 2nd ed. New York, NY: Springer.
- Bowers, K.J., Johnson, S.D., and Hirschfield, A., 2004a. The Measurement of Crime Prevention Intensity and its Impact on Levels of Crime. *British journal of criminology*, 44 (3), 419–440.
- Bowers, K.J., Johnson, S.D., and Pease, K., 2004b. Prospective Hot-Spotting: The Future of Crime Mapping? *British journal of criminology*, 44 (5), 641–658.
- Bradford, B., Jackson, J., and Stanko, E.A., 2009. Contact and confidence: revisiting the impact of public encounters with the police. *Policing and society*, 19 (1), 20–46.
- Braga, A.A., Papachristos, A.V., and Hureau, D.M., 2014. The Effects of Hot Spots Policing on Crime: An Updated Systematic Review and Meta-Analysis. *Justice quarterly*, 31 (4), 633–663.
- Brakatsoulas, S., *et al.*, 2005. On Map-matching Vehicle Tracking Data. *In: Proceedings of the 31st International Conference on Very Large Data Bases*, VLDB '05, Trondheim, Norway. VLDB Endowment, 853–864.
- Brantingham, P.L. and Brantingham, P.J., 1993. Nodes, paths and edges: Considerations on the complexity of crime and the physical environment. *Journal of environmental psychology*, 13 (1), 3–28.
- Brantingham, P.L. and Brantingham, P.J., 1995. Criminality of place. *European journal on criminal policy and research*, 3 (3), 5–26.

- Carriaga, M.L. and Worrall, J.L., 2015. Police levels and crime A systematic review and meta-analysis. *The police journal*, 88 (4), 315–333.
- Cascetta, E., *et al.*, 2002. A model of route perception in urban road networks. *Transportation research part b: Methodological*, 36 (7), 577–592.
- Chown, E., Kaplan, S., and Kortenkamp, D., 1995. Prototypes, location, and associative networks (PLAN): Towards a unified theory of cognitive mapping. *Cognitive science*, 19 (1), 1–51.
- Cohen, L.E. and Felson, M., 1979. Social Change and Crime Rate Trends: A Routine Activity Approach. *American sociological review*, 44 (4), 588–608.
- Couclelis, H., *et al.*, 1987. Exploring the anchor-point hypothesis of spatial cognition. *Journal of environmental psychology*, 7 (2), 99–122.
- Daganzo, C.F. and Sheffi, Y., 1977. On Stochastic Models of Traffic Assignment. *Transportation science*, 11 (3), 253–274.
- Davies, T. and Johnson, S.D., 2014. Examining the Relationship Between Road Structure and Burglary Risk Via Quantitative Network Analysis. *Journal of quantitative criminology*, 31 (3), 481–507.
- Eck, J. E., 1993. The threat of crime displacement. *Criminal justice abstracts*, 25, 527–546.
- Eck, J.E., Clarke, R.V., and Guerette, R.T., 2007. Risky Facilities: Crime Concentration in Homogeneous Sets of Establishments and Facilities. In: G. Farrell, S.D. Johnson and M. Townsley, eds. *Imagination for Crime Prevention: Essays in Honour of Ken Pease*. Crime Prevention Studies, vol. 21. Monsey, NY: Criminal Justice Press, 225–264.
- Famega, C.N., Frank, J., and Mazerolle, L., 2005. Managing Police Patrol Time: The Role of Supervisor Directives. *Justice Quarterly*, 22 (4), 540–559.
- Gärling, T., Laitila, T., and Westin, K., 1998. *Theoretical foundations for travel choice modeling*. Oxford, UK: Pergamon.
- Golledge, R.G., 1999. Human wayfinding and cognitive maps. In: R.G. Golledge, ed. *Wayfinding behavior: Cognitive mapping and other spatial processes*. Baltimore, MD: Johns Hopkins Press, 5–45.
- Hawdon, J.E., Ryan, J., and Griffin, S.P., 2003. Policing Tactics and Perceptions of Police Legitimacy. *Police quarterly*, 6 (4), 469–491.
- Hillier, B., 1996. *Space is the machine*. Cambridge, UK: Cambridge University Press.
- Hillier, B., 2004. Can streets be made safe? *Urban design international*, 9 (1), 31–45.
- Hillier, B. and Hanson, J., 1984. *The Social Logic of Space*. Cambridge, UK: Cambridge University Press.
- HMIC, 1998. *What Price Policing? A Study of Efficiency and Value for Money in the Police Service*. London, UK: Her Majesty’s Inspectorate of Constabulary.
- Jiang, B. and Caramunt, C., 2004. Topological analysis of urban street networks. *Environment and planning b: Planning and design*, 31 (1), 151 – 162.
- Johnson, S.D. and Bowers, K.J., 2010. Permeability and Burglary Risk: Are Cul-de-Sacs Safer? *Journal of quantitative criminology*, 26 (1), 89–111.
- Kelling, G.L., *et al.*, 1974. *The Kansas City Preventive Patrol Experiment: A Summary Report*. Washington, DC: Police Foundation.
- 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.
- Kowalska, K., Shawe-Taylor, J., and Longley, P., 2015. Data-driven Modelling of Police Route Choice. In: *Proceedings of the 23rd GIS Research UK conference*, Leeds, UK.
- Lou, Y., *et al.*, 2009. Map-Matching for Low-Sampling-Rate GPS Trajectories. November, Seattle, WA. ACM.
- Manley, E.J., Addison, J.D., and Cheng, T., 2015. Shortest path or anchor-based route choice: a large-scale empirical analysis of minicab routing in London. *Journal of transport geography*, 43, 123–139.
- Marvell, T.B. and Moody, C.E., 1996. Specification Problems, Police Levels, and Crime Rates. *Criminology*, 34 (4), 609–646.
- Matrofski, S.D., 1990. The Prospects of Change in Police Patrol: A Decade in Review. *Amer-*

- ican journal of police*, 9 (3), 1–79.
- McCabe, J., 2013. *An analysis of police department staffing: How many officers do you really need?* An ICMA Center for Public Safety Management White Paper. Washington, DC: International city/County Management Association.
- Merchant, D.K. and Nemhauser, G.L., 1978. A Model and an Algorithm for the Dynamic Traffic Assignment Problems. *Transportation science*, 12 (3), 183–199.
- Mohler, G.O., *et al.*, 2011. Self-Exciting Point Process Modeling of Crime. *Journal of the american statistical association*, 106 (493), 100–108.
- Newburn, T., ed., 2004. *Policing: Key readings*. Cullompton, UK: Willan.
- Pawson, R. and Tilley, N., 1997. *Realistic evaluation*. London, UK: Sage.
- Pink, O. and Hummel, B., 2008. A statistical approach to map matching using road network geometry, topology and vehicular motion constraints. *In: 11th International IEEE Conference on Intelligent Transportation Systems, 2008. ITSC 2008*, October. 862–867.
- Porta, S., Crucitti, P., and Latora, V., 2006a. The network analysis of urban streets: A dual approach. *Physica a: Statistical mechanics and its applications*, 369 (2), 853–866.
- Porta, S., Crucitti, P., and Latora, V., 2006b. The network analysis of urban streets: A primal approach. *Environment and planning b: Planning and design*, 33 (5), 705 – 725.
- Porta, S., *et al.*, 2012. Street Centrality and the Location of Economic Activities in Barcelona. *Urban studies*, 49 (7), 1471–1488.
- Porta, S., *et al.*, 2009. Street centrality and densities of retail and services in Bologna, Italy. *Environment and planning b: Planning and design*, 36 (3), 450 – 465.
- Prato, C.G., Bekhor, S., and Pronello, C., 2011. Latent variables and route choice behavior. *Transportation*, 39 (2), 299–319.
- Quddus, M.A., Ochieng, W.Y., and Noland, R.B., 2007. Current map-matching algorithms for transport applications: State-of-the art and future research directions. *Transportation research part c: Emerging technologies*, 15 (5), 312–328.
- Ratcliffe, J.H., *et al.*, 2011. The Philadelphia Foot Patrol Experiment: A Randomized Controlled Trial of Police Patrol Effectiveness in Violent Crime Hotspots*. *Criminology*, 49 (3), 795–831.
- Reiner, R., 1992. Policing a Postmodern Society. *The modern law review*, 55 (6), 761–781.
- Reiss, A.J., 1992. Police Organization in the Twentieth Century. *Crime and justice*, 15, 51–97.
- Sampson, R.J., Raudenbush, S.W., and Earls, F., 1997. Neighborhoods and Violent Crime: A Multilevel Study of Collective Efficacy. *Science*, 277 (5328), 918–924.
- Shen, J. and Cheng, T., 2015. Group Behaviour Analysis of London Foot Patrol Police. *In: Proceedings of the 23rd GIS Research UK conference*, Leeds, UK.
- Townsley, M., & Sidebottom, A., 2010. All Offenders Are Equal, but Some Are More Equal Than Others: Variation in Journeys to Crime Between Offenders. *Criminology*, 48 (3), 897–917.
- Wain, N. and Ariel, B., 2014. Tracking of Police Patrol. *Policing*, 8 (3), 274–283.
- Weisburd, D., 2015. The Law of Crime Concentration and the Criminology of Place*. *Criminology*, 53 (2), 133–157.
- Weisburd, D., *et al.*, 2013. The Effects of Increased Police Patrol on Crime and Disorder.
- White, G.F., 1990. Neighborhood permeability and burglary rates. *Justice quarterly*, 7 (1), 57–67.
- Wilcox, P. and Eck, J.E., 2011. Criminology of the unpopular. *Criminology & public policy*, 10 (2), 473–482.
- Williams, S. and Coupe, T., 2017. Frequency Vs. Length of Hot Spots Patrols: a Randomised Controlled Trial. *Cambridge journal of evidence-based policing*, 1–17.
- Willis, J.J., Mastrofski, S.D., and Weisburd, D., 2007. Making Sense of COMPSTAT: A Theory-Based Analysis of Organizational Change in Three Police Departments. *Law & society review*, 41 (1), 147–188.
- Zhao, J.“., Schneider, M., and Thurman, Q., 2002. The effect of police presence on public fear reduction and satisfaction: A review of the literature. *The justice professional*, 15 (3), 273–299.