POTENTIAL USES OF COMPUTATIONAL METHODS IN THE EVALUATION OF CRIME REDUCTION ACTIVITY

by

Shane D. Johnson

University College London

Abstract: *According to Moore's law, increases in computing power are roughly exponential over time. As a consequence, the application of computational methods to old and new problems becomes ever more possible, even with desktop computing. Such methods, and in particular computer simulation, have considerable potential in the study of crime but their application is relatively novel at this time. Consequently, the aim of this chapter is to consider the possibilities with a particular focus on how they might be used to inform the evaluation of crime reduction activity. A number of different types of computational methods will be discussed and examples of the types of policy-related questions for which they might be used considered. The strengths and weaknesses of the approaches described will also be discussed*.

INTRODUCTION

Computational methods, and in particular computer simulation (e.g., Liang et al., 2001; Groff, 2007a), have considerable potential in the study of crime. The aim of this chapter is to discuss some of the possibilities with a particular focus on the evaluation of crime reduction activity. In writing this chapter, two possibilities regarding the scope and direction of the material to be covered suggested themselves: 1) to focus on a particular type of simulation and provide a detailed exposition of that method; or, 2) to discuss a number of different types of computational method to provide the reader with a more general understanding of the possibilities. As the use of simulation for the evaluation of crime reduction activity is novel, I chose the latter. In addition to illustrating the benefits of the methods discussed, some of the issues that need to be addressed before their potential may be realised will also be considered.

The chapter is divided into four sections, each considering different research questions that should inform the various stages of policy decision making. In the first, a concrete example of a technique used to evaluate the effect and sustainability of a situational crime prevention measure implemented at the individual household level is discussed. The method used differs from traditional approaches in that a *Monte Carlo simulation re-sampling* procedure is used to estimate what would have been expected in the absence of intervention and to estimate the statistical significance of observed effects. This type of approach is very different from the simulation methods discussed in the remainder of the chapter, and is used to provide an example of a type of analysis used to answer a very simple “what if” question.

In the second section, discussion moves to the evaluation of area (rather than individual) level interventions. The potential use of *microanalytical-simulation* in the estimation of expected area-level crime rates (in the absence of intervention) is discussed, with a particular focus on how this type of simulation can be used to model data-generating processes not explicitly considered in traditional types of analysis. The aim of this section is to discuss some of the relevant issues rather than to present the results of an actual evaluation. However, an empirical illustration is provided to show how levels of crime may vary in an area even in the absence of intervention under different “what if” conditions.

In the third section, issues regarding the implementation of interventions are considered. Evaluations of place-based crime reduction initiatives generally indicate that implementation is gradual rather than abrupt, and that for successful interventions there is a relationship between the timing and intensity of implementation and the volume of crime prevented (e.g., Bowers et al., 2004). However, rarely are the likely effects of different implementation schedules explicitly considered prior to implementation. Consequently, in this section, using an extension of the methodology discussed in section two, examples are provided of how simulations might be used to estimate the possible effects of different implementation plans prior to the inception of crime reduction activity.

In the final section, the potential use of computer simulation to test theoretical models of crime reduction strategies before they are piloted in an operational context will be discussed. A simple example is provided to illustrate some of the concepts and to provide a focus for discussion.

The material covered in the different sections varies in a number of ways. First, the examples selected consider different units of analysis. In the initial section, the research discussed concerns a micro-level analysis of an intervention implemented at the individual household level. In the second and third, the discussion moves to the evaluation of area-based interventions, and in the final section, the example considered focuses on more general policies that may have no specific geographical boundaries.

Second, the computational methods considered are quite different. I start with a simple example similar to the types of methods with which most readers will be familiar and that can be applied right now. The discussion then moves to methods of simulation that are yet to be used in the evaluation of interventions, or in the testing of crime prevention models but that with a little work could be used at this time. I conclude with a discussion of simulation methods that could possibly be used in the future but that require considerable development before that possibility becomes a reality. To make the chapter as accessible as possible, the use of equations and technical vocabulary is avoided where possible.

1. ESTIMATING THE IMPACT AND SUSTAINABILITY OF INTERVENTION

The first example considers how one might evaluate the impact on domestic burglary risk of a traditional target hardening scheme, where individual households (rather than areas) most at risk of victimisation are targeted for intervention. This particular example is used to illustrate a series of issues associated with this type of evaluation and how the flexibility afforded by computational methods can help reduce threats to internal validity (i.e., rule out other explanations for observed changes).

This type of situational crime prevention intervention usually involves the upgrading of physical security measures at households with identified vulnerabilities (e.g., Forrester et al., 1988). Implementation strategies for this type of intervention vary. For example, individual household surveys may be conducted within high crime neighbourhoods to identify homes with inadequate physical security features which may be vulnerable to victimisation. Alternatively, vulnerability may be identified through an analysis of recorded crime data; research consistently demonstrates that prior victimisation is an excellent predictor of future risk (e.g., Budd, 1999) and thus the prevention of repeat victimisation may represent an efficient burglary reduction strategy (e.g., Farrell, 2005; Pease,1998). In the example discussed below, I consider how one might evaluate an intervention for which the latter strategy is adopted, but this will be discussed only so far as it serves current purposes. A full discussion of the data used and the implications of the findings of this type of evaluation are provided elsewhere (Bowers, Lab and Johnson, 2008) and so only the main points will be covered here.

Approaches to Evaluation

If implemented in a particular area, one approach to evaluation would be to examine changes observed at the area level using an experimental or quasi-experimental design. For both types of design, the counterfactual – what is likely to have happened in the absence of intervention – is estimated by contrasting the crime rates before and after intervention for a treatment and control group, the latter being matched as closely as possible with the former (for an overview of evaluation methods, see Campbell and Stanley, 1963). Where the change observed for the treatment group is preferable to and different from that experienced by the control group by a meaningful amount, the intervention is deemed to have been a success. Some form of statistical significance testing is also typically used to establish whether the size of the effect could have occurred on a chance basis, or is more likely to be attributable to intervention.

Experimental designs allocate target units (people, places or whatever is being studied) randomly to treatment and control groups. This approach is generally preferred because as well as producing matched samples,[1] it eliminates any selection bias[2] that might result from the use of other allocation strategies (Campbell and Stanley, 1963; Shadish et al., 2002). However, when an intervention is implemented in only one or two areas – which is often the case for new types of intervention – (complete) randomization may not produce adequately matched samples. In such cases, either block randomisation[3] or a quasi-experimental design – whereby the control group is selected because of its similarity to the treatment group in as many respects as possible apart from the assignment to the treatment condition (for a further discussion, see the chapter by Henry in this volume) – will be more appropriate.

While this may seem fitting, for an intervention such as target hardening, attempting to estimate the effect of intervention at the area level would be insensitive to the unit of analysis at which implementation occurs. This is so for at least two reasons. First, how does the evaluator define what the geographic area of intervention actually is? One approach would be to use an existing administrative boundary (such as a police beat) which generally encapsulated the area of interest. Even better, a bespoke boundary could be generated using a Geographical Information System (GIS). However, in both cases modifying the boundary used could lead to different results. Referred to as the Modifiable Areal Unit Problem (MAUP; Openshaw, 1984), this is a well documented problem which occurs when data are aggregated at the area level.

---

[1] Matched samples may not be generated using random allocation where the sample sizes are small or where the population is non-homogenous.

[2] Where a selection bias exists, the treatment effect may be confounded with the allocation strategy employed.

[3] Block randomisation is a two-stage process. In the first stage, pairs of candidates (e.g. areas for intervention) are identified that are matched on a range of variables that might influence the dependent variable. In the second stage, one member of each pair is randomly allocated to the treatment condition. For an example in a criminological context, see Braga and Bond, 2008.

Second, unless all homes receive intervention, the evaluation design will be insensitive to expected variation in the change in crime risk to homes that do and do not receive intervention. That is, if measured at the area level, any observable effect of intervention may be diluted as the effect will be measured using data aggregated across two different populations (those that did and did not receive treatment) with different expected outcomes. The extent to which this is a problem will, of course, depend on the dosage of intervention; being more of a problem when relatively few homes receive the intervention. In any event, by not taking account of this, even for an intervention that actually works, the evaluator may underestimate the size of the effect of intervention or make a type II statistical error by assuming that there was no effect where in fact this was simply lost in the aggregation.

Aggregation of the kind discussed also applies to the dimension of time. In the current context, a simple but generally invalid assumption would be that all homes in receipt of intervention were treated on the same day (e.g., the first day of the "after period"). Rarely will this be the case, and implementation may take weeks or even years to complete. However, for the standard before and after design discussed above, the evaluator essentially makes this assumption as crime rates are aggregated for the periods before and after intervention and the effect of intervention is estimated by dividing one by the other.[4] Assuming that the effect of intervention is cumulative, such a design is likely to underestimate any treatment effect.

Additionally, such analysis provides no indication of the longevity of the effect of intervention. Understanding the sustainability of an intervention's effect is important for at least two reasons. First, from a cost effectiveness perspective, the best interventions are those that have a lasting impact on crime. Second, if the probable lifetime of an intervention effect is known, then crime reduction

[4] Time series analyses (e.g. Bowers et al., 2005) do not suffer from this problem. However, such analysis will only be appropriate where data are available for an interval of time which allows adequate diagnosis of ARIMA model parameters and where the volume of crime per unit time satisfies basic requirements. For example, it would be inappropriate to analyse time series data for a small area for which many of the observations (e.g. monthly crime counts) were zero values (a statistical floor effect).

agencies can plan accordingly, timetabling further activity to reinforce intervention.

Survival Analysis

An alternative method would be to see how the risk of crime varies for homes that do and do not receive the treatment before and after intervention. If the intervention is successful then relative to those that do not receive it (the control group), the burglary rate for the treatment group should reduce over time. However, while this approach would deal with the problem of spatial aggregation discussed, it would not address the temporal aggregation problem.

A different approach that could be used, and one that is used frequently in research concerned with recidivism (e.g., Visher and Linster, 1990), is survival analysis. For this type of analysis, the question of interest concerns the typical time to failure, however defined. In the context of studies of recidivism, this is the elapsed time between the start of a rehabilitation programme (for example) and the first offence committed following the start of treatment. For a situational crime prevention intervention, this would be the elapsed time between the treatment and the first victimisation post-intervention. A problem may arise with this type of analysis when the evaluation period covers an insufficient interval of time to allow the time to failure to be identified for all experimental units. For example, consider a household that received an intervention just before the end of the evaluation period. In this case, the likelihood of estimating the true survival time for that home would be low as the data set would essentially be truncated. However, as this problem of what is known as *censoring* is well understood (Tabachnick and Fidell, 2001) and can be corrected for it will be discussed no further here.

The usual approach taken to analysis is to compute a survival distribution (presented as a curve) for those in receipt of treatment and to compare this with those for a suitably matched control group, controlling for the problem of censoring discussed above. In the case of offender rehabilitation, assignment to conditions can often (though not always) be achieved using random allocation. For other types of intervention, random allocation to treatment and control groups is not always possible or even appropriate (for a general discussion, see Sherman, 2007), and hence an alternative matching procedure will be required.

Despite its popularity in the evaluation of medical treatments and offender based programs, hitherto survival analysis has not been used for the evaluation of situational crime prevention measures and so an example of how this method could be used for this type of intervention will be presented along with a discussion of how a suitable control group might be constructed when random allocation is not possible.

Monte Carlo Simulation

As discussed, in the case of survival analysis, the approach to analysis is very similar to the experimental logic discussed at the beginning of this section. Survival curves, which show the cumulative percentage of households unvictimised per unit of elapsed time, are computed for both the treatment and control groups and the patterns compared. To identify a control group, the basic approach would be to describe the treatment group and identify a control group which comprises a set of homes with similar characteristics.

However, for the example discussed here it is possible that for some homes subject to intervention, there will be more than one household that could serve as a control, meaning that the sample sizes would vary across groups if all possible controls were used in the analysis. This may not be problem and one could include all of the potential control households in the analysis. However, as there may be more potential control households for some homes than others, this could lead to a control group for which the overall profile is quite dissimilar to that for the treatment group. The implications of such an aggregation effect for causal inference should require no articulation.

An alternative approach would be to pick a random sample of control households from those available, one for each home in receipt of treatment. The survival curves for the two groups could then be compared and the statistical significance of any differences tested in the usual way. However, the problem with this approach is that any differences observed could be due to the control sample selected: a control sample selection effect.

A different solution is to use a Monte Carlo (MC) re-sampling technique. Here, rather than selecting one control group, a sample of control groups – with each home matched with a treatment household on a one-to-one basis – may be drawn from all permutations possible. For each sample selected, a comparison can be made with the results for the treatment group. Where a large number of samples are drawn (say 99), the average survival curve may be computed along with the standard error. The statistical significance of any observed differences can also be easily estimated, either using the standard errors computed, or more directly by counting for how many of the control samples the survival curve exceeds that for the treatment group (see North, 2002).

An Empirical Example

To illustrate the approach in a more concrete way, the results presented in a recent paper are summarised (Bowers et al., 2008) here. In that study, the effectiveness of the target hardening of individual households was examined for one area on Merseyside (U.K.). The treatment group consisted of 318 households for which intervention had been triggered as a consequence of victimisation experience.

The data available for analysis included the following:

a) The location of all homes within the area that had received target hardening (see Bowers et al., 2008);
b) recorded burglary data for a six year interval; and,
c) Ordnance Survey (OS) data that indicated the address points of every house in the area.

Software was written to identify every home within the area (of which there were 71,500) that did not receive the intervention but that had a similar victimisation profile to those that did. In so doing, the potential problem of regression to the mean (RTM) was reduced (see Campbell and Stanley, 1963). To elaborate, the problem of RTM would occur if homes were assigned to the treatment condition on the basis of an extreme score on a variable (e.g., the number of victimisations experienced in the recent past) that was atypical for that household. The issue is that even in the absence of intervention, over time the risk to such homes would be expected to regress back to the “normal” level generally experienced by them. Where such homes are allocated to the treatment group this may create the illusion of a treatment effect. However, for the current methodology, as homes were matched on the very variables that could lead to RTM being mistaken for something more than it is, the patterns observed for both groups are equally likely to be caused by RTM. Thus, any difference observed between groups should be attributable to an intervention effect (or at the very least something other than RTM).

On average, for each treatment household, there were around 97 homes that fitted the criteria applied, meaning that a large number of control groups could be identified. A MC simulation was used to (re)sample from the universe of control groups possible, and a distribution of survival curves generated. Figure 1 shows the results of the analysis and indicates that relative to the control group(s), homes that received target hardening were typically less likely to be victimised for a period of around two years. The dotted lines shown in Figure 1 show the 5th and 95th percentiles of the MC simulation. Where the observed values exceed the upper dotted line, this indicates that the treatment group exceeded chance expectation for that point in time. The maximum value on the x-axis is 150 weeks which represents the maximum interval of time for which any of the households were protected by target hardening for the period of time for which data were available.

For the current analysis, the results show that the difference in the survival curves was statistically significant for a period of about 18 months. Thereafter, the proportion of homes victimised was roughly equivalent between groups. Thus, for these data at least, target hardening appears to have an immediate crime reductive effect that was sustained for around two years.

INSERT FIGURE 1 ABOUT HERE

Observations and Limitations

Re-sampling from all households using a computational approach and recorded crime data has the clear advantage that unlike other research designs that use surveys, the time frame for analysis need not be limited; analyses can be conducted using as many or as few homes as necessary; construction of the control sample is inexpensive; and, the criteria used to identify the control group(s) can be varied in as many ways as the data permit. Thus, the researcher can ask whether the result obtained is likely to be statistically significant for control groups configured in different ways. In the current example, only simple criteria were used but other specifications are possible. Such flexibility would be impossible to achieve without intensive computation for anything other than very small sample sizes, and data collection would be expensive where surveys were required.

However, it is important to note that the approach will only be as good as the data available. For example, one disadvantage with the data used above is that only limited information was available about each home. For those victimised, information may be available regarding the type of home, who owns it (and so on), but such data are unlikely to be found in police data for the remainder of the population. This may pose a problem if one wishes to match households (treatment and controls) on a range of characteristics that may be observable to the would-be offender.

It is possible that this issue may be minimised by using alternative data such as those collected as part of a government census. For example, in the UK, although unavailable at the individual household level, data are available at a fairly high level of resolution, the most precise being the census Output Area (OA). Each OA contains around 125 homes, and areas are defined to maximise within area homogeneity; that is, as far as possible the geographical boundaries derived delineate areas that maximise the likelihood that similar people live within the same OAs. Thus, although this would not completely resolve the issue, using the OA geography it would be possible to estimate the probability that any particular household shared characteristics with those with which they are to be matched.

As further issue that re-sampling does not entirely resolve is that of selection bias and the related omitted variable problem. For example, in the current example it is possible that there was something systematically different about those homes that did and did not receive treatment, but that this was not apparent from the available data. It may be that this difference explains the variation in the survival curves. This is difficult to overcome in the absence of random allocation or where the data required to more precisely match the two groups – treatment and control – are unavailable. Where more extensive data are available it would be possible to refine the procedure using the propensity score matching approach to sample construction. The aim of the approach would be to ensure that the homes in the two conditions would both be equally likely to have been assigned to the treatment group based on the inclusion criteria adopted (for more details, see the chapter by Henry in this volume). In this case, any observed differences between groups could be attributed to the intervention with increased confidence.

For these reasons, the results generated by techniques such as that described should be interpreted in a sensitive way and used in studies that are carefully designed to minimise as many threats to internal validity as possible. However, it is also important to recapitulate the advantage of using a re-sampling methodology and how this can help to minimise bias. First, as the construction of the control groups is iterative, any observed effects are unlikely to be attributable to peculiarities in the samples identified. Second, as this approach to hypothesis testing enables p-values to be calculated directly using the re-sampling methodology, the data need not conform to a particular statistical distribution, which is a requirement of most statistical tests. Finally, as the approach essentially involves the use of many control groups, there is no risk of making errors of inference that could arise from using only one control group. Consider that if one control group were selected from all those possible, the conclusions drawn would be affected by this selection. If the control group used represented an extreme case (e.g., if it was below the 5th percentile of MC Simulation) and standard statistical tests were used this would lead to an error of inference.

2. ESTIMATING THE IMPACT OF INTERVENTION AT THE AREA LEVEL

There are at least three types of policy question for which computer simulation may be useful. First, simulation may be used to estimate what patterns of crime would be expected in a given area (or for a given population) in the absence of intervention. These estimates may then be compared to patterns observed to determine if an intervention is likely to have had any effect. Second, simulation may be used as a tool for testing the likely effects of an intervention for a range of implementation scenarios in which the timing and intensity of activity is varied. Third, simulation may be used to systematically test theoretical models of interventions before expensive field trials are conducted. In this section the first possibility will be discussed, while the others will be considered in subsequent sections.

Before continuing it is worth outlining some of the main differences between simulation as method and more traditional approaches. Many analytic approaches to theory testing employ a top-down methodology; patterns are observed in the real world and the data generating processes or mechanisms for them inferred. Using simulation, a bottom-up approach is adopted. That is, a data generating process is specified *a-priori* and the simulated phenomena that emerge (e.g., simulated patterns of crime) observed. Put another way, a computer simulation is an implementation of a theory. Thus, much like a thought experiment, one can ask simple “what if” questions under conditions where variables of interest can be manipulated and their effects – along with those of chance – assessed.

A range of simulation methodologies exist (for a general review, see Gilbert and Troitzsch, 2003) and some of these have already been applied to the study of crime (for a review, see Alimadad et al., 2008; and for a collection of examples, Liu and Eck, 2008; Groff and Mazzerole, 2008). For example, McAllister et al. (1991) use a Queing simulation model to examine case processing within the court system in the US and estimate the impact of policy changes on the efficiency of the system. Johnson (2008) uses microanalyitcal-simulation to test theories of crime concentration. A number of researchers (e.g., Groff, 2007a; Birks et al., 2008) have used agent-based models to examine routine activity theory (Cohen and Felson, 1979) and the effects of police patrols on offender activity (e.g., Dray et al., 2008). However, before getting too excited about advanced methods of simulation, I will discuss a simple example that uses microanalytical-simulation.

INSERT BOX 1 ABOUT HERE

Estimating the Counterfactual

Techniques for estimating the statistical significance of experimental manipulations conducted in the laboratory (or medical trials) are quite simple. In the simplest case, a group of people may be randomly assigned to one of two different conditions (experimental and control), and their performance on a test observed. As already discussed, random allocation to conditions will generally lead to the formation of two groups with similar characteristics. This minimises the problem of selection bias, meaning that the results obtained are unlikely to be due to the way in which participants were allocated to conditions.

The types of statistical tests used to determine whether between group differences observed are meaningful generally involve the estimation and consideration of the standard error of the sample means. The larger the sample sizes, the more reliable the estimates. The difficulties of using this paradigm and these types of statistical test in studies of crime prevention implemented at the *area*- (rather than person or household-) level are numerous, but three of the main points are listed below:

- The randomization of allocation of areas to conditions will rarely be possible or, given the highly skewed distribution of crime risk (e.g., Sherman, 2007) and the variation in context across locations (e.g., Tilley, 1993), even appropriate.

- Many studies involve only **one** treatment and control group and these will often be implemented across geographic areas (e.g., studies of street lighting) rather than targeting particular individuals or buildings.

- And, for problem-oriented policing projects, the crime reduction strategy should be designed to interfere with the specific conditions that facilitate crime in the area selected for intervention. It is assumed that these conditions will vary across locations and so the same solution would not be expected to work across all or any conditions. Thus, for problem-oriented projects, the random allocation of areas to conditions would be incongruent with the philosophy of the approach.

A number of alternative approaches have thus been used for the evaluation of area-based interventions (Ekblom and Pease, 1995; Johnson et al., 2004; for a classic overview, see Campbell and Stanley, 1963). For most, the basic assumption is that the crime rate in an area post-intervention will be a function of the crime rate before intervention multiplied by a coefficient of expected change plus some degree of chance fluctuation. The expected change is usually estimated by examining observed changes in similar areas (e.g., Farrington and Welsh, 2006), or the wider area within which the intervention area is located. The statistical significance of a result obtained is then estimated by computing a standard error for the estimate derived. Thus, where a change is observed in an intervention area, the evaluator attempts to rule out threats to internal validity (alternative explanations) for the change by seeing how things changed in similar areas not subject to intervention. Where the change in the intervention area exceeds those in other areas by a substantial amount, the change observed *may* be attributed to the intervention with more confidence.

In addition, if the evaluator wants to understand how the intervention may (or may not) have worked, and to strengthen conclusion validity, he will usually want to identify the likely crime reduction mechanisms (see Pawson and Tilley, 1997) through which the intervention could have worked and, through the collection of the relevant data regarding intermediate outcomes (e.g., change in the number of residents that have noticed new street lighting), see if these mechanisms were triggered. Where they are not, causal inferences will be challenged or at the very least the mechanisms through which the intervention may have had an impact will need to be revised. The importance of conducting such hypothesis testing and identifying signatures that bespeak mechanism is difficult to overstate (for a detailed discussion, see Pawson and Tilley, 1997 and the chapter by Eck and Madenson in this volume), but for reasons of space will not be discussed further here.

Returning to the estimation of the impact of intervention, using the basic logic discussed above, one simple statistic that may be calculated is the odds ratio (see Farrington and Welsh, 2004). This is computed by comparing the ratio of change (before versus after) in an intervention area to that in a comparator. An odds ratio of 1 so derived would indicate that the changes observed were equivalent across areas. A value above one would indicate a reduction in the treatment area relative to that observed in the comparator. The standard error of the estimate is computed by assuming that the number of crimes in an area conforms to a Poisson distribution, and hence that the variance will be equal to the mean. A potential problem with this approach is that Poisson models are susceptible to the problem of overdispersion whereby estimates so computed may underestimate the variability observed in the real world (e.g., see Agresti, 2002).

Reasons for simple Poisson models underestimating variability relate to the fact that they do not model all of the processes likely to affect the dependent variable. The assumption of independence of events is potentially problematic (see Marchant, 2005; but see Farrington and Welsh, 2004). For example, research demonstrates a dependency in the timing and location of crime events such as burglary (Townsley et al.,2003; Johnson and Bowers, 2004ab; Bowers and Johnson, 2005; Johnson et al., 2007), vehicle crime (Johnson et al., 2008) and gun crime (Ratcliffe and Rengert, 2007). Theoretical (e.g., Johnson and Bowers, 2004b) and empirical work (Bernasco, 2008; Johnson et al., 2008) suggests that the reason crimes cluster in space *and* time in this way, and do so more than would be expected on the basis of area-level variation in risk, is that the patterns observed reflect the use of optimal foraging strategies by offenders. Simply put, analyses of crimes detected by the police (Bernasco, 2008; Johnson et al., 2008) demonstrate that having victimized one home, offenders often target the same home again and others nearby within a short space of time. There is dependency in the data.

A potential advantage of microanalytical-simulation over standard statistical methods is that space-time processes (where they are known to the researcher and can be formally expressed) can be incorporated into the models enabling their effects on the dependent variable to be estimated. The need to do this was actually discussed some time ago by Barr and Pease (1992) who highlighted the importance of understanding and taking into account normal patterns of crime placement in evaluation research. However, as far as I am aware, such models have not been developed or discussed in any detail hitherto.

An Earlier Experiment Conducted Using Microanalytical-Simulation

To illustrate how this might be done, an example of a microanalytical-simulation model will be discussed and the results of a simulation experiment presented. The model was originally developed (see Johnson, 2008) for the purposes of theory testing and falsification, with a focus on patterns of repeat burglary victimisation. Over three decades of research demonstrate that prior victimisation is an excellent predictor of future risk for all crime types so far studied (for reviews, see Pease, 1998; Farrell, 2005). However, while the ubiquity of such findings is largely uncontested, debate still exists regarding the theoretical mechanisms that generate the phenomenon. Some argue that patterns of repeat victimisation may be explained by enduring heterogeneity of crime risk across homes; some homes are simply more vulnerable than others and consequently are repeatedly victimised. Others (e.g., Pease, 1998) suggest that a first offence increases the probability of future victimisation, either because the victim’s behaviour changes in response to the first offence, or because the offender's newly acquired knowledge of that home increases the attractiveness of it relative to those that remain unknown.

I will not rehearse the findings of the substantial research literature concerned with repeat victimisation any further, but instead discuss how the model was developed to test the two theories, and how it could be used for the purposes of estimating the counterfactual. In the earlier study, patterns of crime were simulated for a virtual population of homes under a range of conditions. The central question was whether patterns of repeat victimization as observed in the real world could be generated by a microanalytical-simulation where the risk to individual homes was the product of time-stable risk factors and the roll of a virtual dice, or whether a further mechanism was required.

To test the first question, every home in the simulation, generated using point data for a whole county in the U.K., was allocated a particular risk of victimization. The latter was estimated using police recorded burglary data for a period of four years. Over a simulated interval of four years, on each virtual day some homes were selected as burglary victims while others were not. The selection of which homes were victimized each day was determined by the risk allocated to each home and the output of a uniform random number generator (RNG). Homes with the highest risks were selected more frequently than those with the lowest, but the model was not deterministic and as would be expected the results varied across runs.

In addition to modeling time-stable risks, an element of what I will refer to (for the sake of simplicity) as "contagion"[5] was included in some of the models whereby the probability of victimization at each home increased as a function of the (virtual) burglaries experienced at that home. In addition to modeling these factors, seasonality was included in the models by varying the risk of crime according to the variation observed in the police recorded crime data.

[5] There is no suggestion here that changes in the risk of victimisation are likely to be the consequence of a biological or similar mechanism. The simile is merely useful as a mental shortcut.

More precise details of the model and the procedures used will not be discussed here (see, Johnson, 2008), but it is worth noting that the risk to each home was estimated using the police recorded crime data and the U.K.OA census geography discussed above. Other methods of deriving the estimates of risk exist, and the use of a particular method may, of course, affect any results generated. In this case, the most significant threat to ecological validity is the Modifiable Areal Unit Problem (Openshaw, 1984), whereby changing the boundaries used in the analysis may affect the estimates derived (for a discussion of other issues, see Johnson, 2008).

Model Validation

An important aspect of simulation research concerns model validation. Townsley and Johnson (2008) discuss a framework for analysis which draws upon the validity typology conceived by Campbell (1957). A lengthy discussion will be avoided here, but some of the questions to be asked include whether constructs formalized as part of the model accurately reflect the intended definition; if threats to internal validity could have evolved from coding or other errors; and whether the results approximate those in the real world (empirical validity). Threats to validity are likely to increase with the complexity of the simulation.
In the case of the Johnson (2008) study, the model was very simple and so only the results generated will be discussed. Simply put, these suggested that whilst area-level crime rates explain some of the variance in crime concentration at the household level, they far from exhaust it. Indeed, only when an element of contagion was introduced did the models generate patterns of victimization that resemble those to be found in police data. A particularly important finding was that time-stable population heterogeneity failed to explain the ubiquitous time course of repeat victimization whereby the risk of revictimization decays exponentially with elapsed time (e.g., Polvi et al., 1991).

For those particularly interested in this type of simulation and how varying the model parameters may affect the results, a variant of it was produced for this chapter. The model was developed in the NetLogo programming language (Wilensky, 1999) which is a cross-platform multi-agent programmable modeling environment which readers can download (http://ccl.northwestern.edu/netlogo/) free of charge.[6]

Using Microanalytical-Simulation to Estimate the Counterfactual

If computer simulations can be developed that represent reasonable approximations of how patterns of crime vary in the absence of particular interventions, then such models could be used to estimate the counterfactual. To illustrate, the model described above was here used to explore the potential effects of the space-time dependency of crime placement on the estimation of the counterfactual. To do this, a geographic area was selected from the virtual world at random, and estimates of the volume of crime expected in it generated for a fictional one year period. The area selected, shown as Figure 2, had a total of 5,583 homes within it.

INSERT FIGURE 2 ABOUT HERE

Estimates of the expected volume of crime in the selected area and the associated variance were computed for three conditions, as follows:

1) Where the occurrence of crime at each home depended only on an estimate of time-stable risk at the neighbourhood level for a notional pre-intervention period. The period used to calibrate the model was the same four-year interval used by Johnson (2008).

2) As per model #1, but with a contagious repeat victimisation process; the risk of victimisation at each victimised home was elevated by a factor of 5 for a period of 1-16 weeks (selected using a uniform RNG). The contribution of time-stable area-level variation in risk was scaled to ensure that the mean count was equal to that for model #1.

[6] To use the model, the reader should first open the NetLogo program. Once running, the user can load the simulation model using the File command. Details of what the simulation does and how to use it can be found in the "information" tab. The simulation model may be downloaded from: http://www.jdi.ucl.ac.uk/british_academy_network/history/index.php

3) And, as per model #2, but with the inclusion of a more general space-time process, whereby the risk to homes within 50 meters of previously victimised locations was temporarily elevated by a factor of 1.1 and for those within 50-100m, by 1.05. As with model 2) the elevation in risk was temporary.

For each model, the simulation was run 100 times. For the purposes of illustration, hotspot maps produced from the data generated by the first two realisations of model #3 (for the wider area within which the study area was located) are shown in Figure 3. It should be evident that similarities exist between the two maps but the precise patterns vary. As the patterns are generated for a fictional period, differences in the patterns observed reflect only the effects of the data generating processes specified, nothing more.

INSERT FIGURE 3 ABOUT HERE

Table 1 shows the average volume of crime for the study area for each model. The results indicate that the variance of the estimates is a function of the processes modelled. This suggests that modelling the factors which influence crime placement is important in the evaluation of crime prevention schemes. Failing to do so will mean that the expected variance will be underestimated and errors of inference will be more likely. A further advantage of using simple simulation models is that where hypotheses are tested, p-values can be calculated directly without the use of statistical tables, meaning that there is no requirement that the data fit a particular distribution (see above).

INSERT TABLE 1 ABOUT HERE

In addition to estimating the direct effect of intervention, a simulation of this kind could be used to estimate the extent of any geographic displacement or diffusion of benefit (see Eck, 1993) that may have occurred. The advantage of so doing is that where Space-Time processes are modeled, these would be less likely to be mistaken for spatial displacement, or target switch.

A final benefit of using simulation in this context is that the possible influences of other changes that may occur at the area-level can be modelled. For example, it is possible that during an intervention period the population at risk may change. New homes may be built or old ones demolished. This too can be modelled by adding or removing homes from the simulation. It is also possible that the residential population may change in a way that might be expected to affect area-level crime rates. For instance, over the period of intervention, the population of residents aged 13-17 years – the peak age in offending (Farrington et al., 2006) – may increase. The model could be calibrated to estimate the effect of this, but with the understanding that results obtained would be dependent upon the assumptions made. Other possibilities exist.

Before getting too excited, a note of caution is necessary. In the above examples, models (2) and (3) were a combination of a mixed Poisson and space-time process. Before these types of model can be used for the purposes of evaluation, we need to better understand the influence of the two processes on area- (and individual household-) level crime rates. What should the contribution of each process be? Does this vary by area? How do we calibrate the models? These are questions that require answers. At the very least, the use of the simulation model highlights the need to identify the relevant parameter values and how these might be modeled and validated. Again this requirement of specificity is a benefit of simulation in that it helps to inform the agenda for basic research and understanding.

It is worth noting that for the models tested, the difference in the variance observed across the models tested is perhaps not as dramatic as one might expect. It may be that this is will not be true for other models in which the Space-Time (or other) processes are more accentuated (or other factors modelled), but one interpretation of this finding is that simple approaches such as the odds ratio method (particularly where conservative estimates are used, e.g., Farrington and Welsh, 2006) may not be as problematic as has been suggested (Marchant, 2005). It is beyond the scope (and aim) of this chapter to explore this further, but this could be investigated more systematically using simulation models.

3. EXAMINING THE LIKELY EFFECTS OF DIFFERENT IMPLEMENTATION STRATEGIES

A second type of question that the above simulation may be useful in helping to address concerns the *dosage* of an intervention and the timing of implementation required to deliver a desired effect. Consider a scenario where an intervention has been selected for implementation and there is a desire to achieve a particular reduction in crime by a given date. Simulation might be used to help determine whether this degree of reduction is plausible and to determine what model of implementation would be required to achieve it.

For a burglary reduction intervention, some of the factors that will influence the impact on crime are the effect of the intervention at the individual household level and the dosage of implementation; the more homes that are treated the greater the expected effect (e.g., Ekblom et al., 1996). Simulation may be used to model different scenarios for a variety of assumptions. For example, if we believe that an intervention will reduce the risk of crime by 20% (or within some range of it) for those in receipt of treatment then this effect can be modelled for different levels of implementation, and the effect on area-level crime rates observed.

Implementation can vary in ways other than dosage. For example, homes could be selected for intervention randomly, or those with the highest risks could be given priority; implementation could be done relatively abruptly within a short space of time, or it could be more gradual, taking months or even years to complete. Simulations could be used to examine the potential impact on crime of a range of models of implementation. Assumptions would need to be made, but across a series of runs, parameters of the model could be varied to examine the effect of intervention for different sets of assumptions.

To illustrate the potential use of this method, two simulation experiments were conducted. In the first, the dosage of intervention was varied and, in the second, the timing of intervention considered. For both experiments, as the aim was one of illustration alone, the simplest model of crime placement discussed above (model 1) was used as a baseline model.

Modeling the Effect of Varying the Dosage of Intervention

In the first simulation experiment, different fractions of homes located within the area used above were selected for intervention on the first day of the simulation. Two different levels of dosage are considered: 50 and 100%. Where the dosage is 50%, two different models of allocation were used to identify homes that received the “intervention.” In the first, homes with the highest “pre-intervention” risk were selected. In the second, those with the lowest risks were chosen.

At the individual household level, for the sake of simplicity the intervention is assumed to reduce the risk of crime to those selected by a factor of 50%. Varying this parameter will, of course, affect the results considerably and it is possible that in the real world the effect of intervention (or even the precise intervention implemented) will vary across homes. However, for simplicity a fixed effect is assumed.

Table 2 shows the results observed across 20 runs of the simulation. The first column shows the average effect of changing the risk to a given fraction of homes by 50%. This is derived by dividing the average count of crime in the treatment area for the intervention model by that for the baseline model (for which the effect of intervention is excluded). The second column shows the worst case scenario. To calculate this, the maximum count of crime generated for the intervention model is divided by the mean for the baseline model. This is not strictly the worst case scenario of course.[7] The final column shows the best case scenario for which the lowest count of crime for the intervention model is compared with the mean count of crime for the baseline model.

Thus, where 100% of homes are selected for intervention this reduces the risk to them by 50%. The results of the simulation suggest that average effect of

---

[7] We could divide the highest count of crime for the treatment area by the lowest count for the baseline model. However, the mean represents what would be typically expected in the absence of intervention and is thus the denominator used here. For the estimation of the treatment effect, in reality time only flows one way and so we will observe only one outcome. For this reason, I use the highest and lowest estimates in this case (there will actually be no average in reality).

so doing is a reduction in crime of around 49%.[8] However, when we take account of random fluctuation (excluding Space-Time processes), the results suggest that the effect measured could actually vary between 38-63%. Thus, on the basis of the simulation results, it is plausible that in the real world if the risk to homes was actually reduced by 50%, the effect observed could be as low as 38%, or as high as 63%.

The minimum and maximum effect sizes vary considerably, particularly when only 50% of homes are selected for intervention. This is due to the fact that there is considerable variation in risk across the area. Thus, if a practitioner were to assume that randomly selecting homes for intervention was as good a strategy as any, he would commit the ecological fallacy (e.g., Bowers et al., 2005) by assuming that all homes within the area experience the same risk of victimization. If the assumptions of the model tested are correct, the outcome of such a strategy would be a smaller reduction than could be achieved, and in the worst case, no effect at all.

INSERT TABLE 2 ABOUT HERE

Estimating the Effect of Using Different Implementation Timelines

In the second experiment, the dosage of intervention is held constant at 50% across simulations, but the timing of intervention varied. For the first model, the simulated implementation phase was abrupt with all homes being treated on day 1. The results for this model are the same as those shown in row 1 of Table 2. For the second model, the rate of implementation was decreased to 17 homes per day, which generated an implementation period of one year.

Unsurprisingly, the results shown in Table 2 indicate that a smaller crime reductive effect is likely to be observed in the one year period when a gradual implementation schedule is used. Of course, the value of the simulation approach is that it can provide an idea of exactly what the effect of intervention might be for a range of different implementation schedules (and assumptions). A

---

[8] This is an estimate of the mean effect, generated over 20 simulations. The true value of 50% would be observed if a larger number of simulation runs was used.

simulation of this kind could also be used to estimate how many homes would need to be treated if an effect of X% was required, but only one particular implementation schedule possible (e.g., gradual implementation).

This general approach could also be used to try to better understand the impact of an intervention where an observable effect is produced. For example, if an intervention is believed to reduce crime, and the timing and dosage of intervention are known, then the parameters of a simulation model could be tuned using this information and the other parameters of the model (e.g., the size and sustainability of the effect, possible diffusion of benefit and so on) varied until the model produced results similar to those observed at the area level in the real world. If the results of the model approximate (or do not) those observed, this could help the evaluator test a range of hypotheses.

Additionally, simulation models might be produced to generate “signatures” that bespeak mechanism so that the emergent patterns could be looked for in real world data. For example, instead of assuming that the effect of intervention is permanent, models could be produced to estimate the effect of intervention where the effect decays over time. Models could be produced to show the impact of intervention where a diffusion of benefit occurs, reducing the risk of crime to homes nearby not subject to treatment.

One obvious limitation of the approach described is that it applies to crimes committed at individual properties. Consequently, the application of the approach will be meaningful for some crimes such as burglary (domestic and commercial), shoplifting, and possibly criminal damage, but may be of less utility for crimes such as street robbery. For the latter it is possible that the same approach could be used if a different unit of analysis were employed. For example, in an analysis of police recorded crime data which included a range of crime types, Weisburd et al. (2004) illustrate the value of using street segments as the unit of analysis. Thus, instead of using homes, other simulations could use street segments.

A further limitation of the approach is that the results are only of value if the assumptions made are valid. This is a general issue with evaluation research where the impact of intervention is to be estimated, but something to which I will return in the next section.

4. SIMULATION FOR POLICY THEORY TESTING

.... consider the scenario of a child at the beach letting sand trickle down to form a pile. In the beginning, the pile is flat, and the individual grains remain close to where they land. Their motion can be understood in terms of their physical properties. As the process continues, the pile becomes steeper, and there will be little sand slides. As time goes on, the sand slides become bigger and bigger. Eventually, some of the sand slides may even span all or most of the pile. At some point, the system is far out of balance, and its behavior can no longer be understood in terms of the behavior of the individual grains. The avalanches form a dynamic of their own, which can be understood only from a holistic description of the properties of the entire pile rather than a reductionist description of the individual grains: the sandpile is a complex system.

(Per Bak, 1997, p. 2).

Some of the complexity associated with patterns of crime was alluded to in a previous section in our discussion of offender as forager (see also, the chapter by Ekblom in this volume). In this section, the ecology[9] of crime (for an extended discussion, see Felson, 2005) is considered in more detail and one form of simulation method that may be used to capture the associated complexity and how this might inform policy decisions discussed.

When thinking about solutions to crime problems, the general approach is to consider the problem as currently conceived and attempt to identify points for intervention that might interrupt the cycle. Ideas may focus on different approaches including forms of offender rehabilitation, sentencing policy, changes in the number of police on the street, or situational crime prevention measures. Whatever the approach considered, those involved in making decisions typically do so by reviewing the available evidence (where it exists), and contemplating

[9] Although the use of simulation is relatively new in the field of criminology, it has been used for decades in the field of ecology to study behaviour not dissimilar to that discussed here. For example, Pyke (1981) describes a computer simulation used to test theories of optimal foraging in Honeyeaters.

the crime reduction mechanisms (Pawson and Tilley, 1997) through which a particular intervention might work. Where a particular intervention is likely to have a fairly simple crime reduction mechanism, this approach may work well, helping policy makers to sift through contending approaches to crime reduction.

However, where the effect of an intervention may be influenced by the reactions of a variety of interacting actors (e.g., offenders, police, and victims) thought experiments may be insufficient. Moreover, because the behaviour of the system may not be understood in reductionist terms, thought experiments may be prone to error. To illustrate, consider the example of police patrols. Changes in the location and timing of police patrols may affect offender targeting choices; offenders may prefer to avoid areas when police patrols are visible (or anticipated), or they may simply wait for patrols to leave before resuming their activities. Whatever their decisions, the actions of offenders may influence the plans made by the police and vice versa. This process is iterative and the patterns that emerge may quickly become too complex for a simple thought experiment. Now imagine what happens when other actors who might influence the probability of crime occurrence are considered, and what happens when we acknowledge that the different actors (offenders, victims and guardians) may not represent mutually exclusive populations.

The modelling of complexity of this kind requires a different kind of simulation than has been so far discussed. *Agent based* methods generally describe simulations in which “agents” are represented by self-contained programs that control how they perceive their “world”, what actions they will take and how they interact with their environment and each other (for a discussion, see Gilbert and Troitzsch, 2003). Agents are usually autonomous so that their “decisions” are not directed by others, but they may communicate with other agents and their interactions with them may influence their behaviour. Different *classes* of agents can be used to represent a range of populations (e.g., offenders and victims), and models can easily be programmed to incorporate population heterogeneity (e.g., offenders with different rates of offending, victims with different risks).

INSERT BOX 2 ABOUT HERE

Returning to our crime prevention example, rather than simply trying to think through what might happen under a given set of conditions (which will require considerable mental gymnastics), agent-based simulations may be implemented to systematically test possible outcomes. Simulations may be run many times to see if the results vary as a consequence of chance effects or starting conditions, and to help identify the conditions under which and the mechanisms through which (optimum) impacts are most probable, and where little or no impact is likely to be observed, or where intervention might lead to unintended consequences (positive or negative).

Of course, to do this the simulator needs to have a fairly good model of how the world works in the absence of intervention (but so too does the practitioner who relies on the thought experiment). This illustrates an important possible limitation of simulation at this time; the utility of the simulation will be a function of how well the model is specified and the extent to which this reflects the way the world might be. Ecological theories of crime (Cohen and Felson, 1979; Ekblom, 2000) provide a useful framework for analysis but, as will be discussed below, more empirical work is required before simulation models can be sufficiently specified.

Nevertheless, simulation offers the potential to improve upon the thought experiment considerably, and to generate results concerning what an intervention might achieve, given a set of clearly specified assumptions. Rather than a weakness this may be considered a further benefit of simulation. That is, the requirement to precisely specify a computer model, which incorporates the decision making processes of all agents, requires the researcher to carefully define the theoretical model used, and to specify the crime reduction mechanisms or logic of the intervention, something that they should always do.

The algorithms used in agent based models are mathematical formalisms intended to mimic or “model” human decision making processes. The basic idea of studying offender decision making is not new. It is central to rational choice theory (Clarke and Cornish, 1985) and has been explored in detail in the context of crime scripts (e.g., Cornish and Clarke, 2003; Lacoste and Tremblay, 2003; see also, Brantingham and Brantingham, 1993). In the case of the latter, the generation of crime scripts involves the identification of templates of behaviour that describe the individual actions (and their sequence) of the various actors (offenders, victims, guardians and so on) involved in crime events. The difference between the approaches is that agent based models can be used to see what patterns emerge when these types of model are implemented in-silico. The generation of crime scripts is a top down analysis, whereas agent based simulation is a bottom-up form of analysis for which crime scripts, or their mathematical analogues define a data generating process.

It is important to note that the algorithms used in agent based simulation are not intended to represent a general theory of cognition, but are instead used to focus on a limited set of decision making processes thought to be important in the behaviour to be understood. While it has limitations, agent based simulation offers the potential to examine models of crime causation and prevention in a more systematic and flexible way than has been possible hitherto.

An Example of an Agent Based Simulation Model

To illustrate some of the issues with agent based simulation, both positive and negative, a concrete example is provided. The example concerns the potential impact of different police patrol strategies on the crime of burglary (see also Birks et al., 2008; Groff and Birks, 2008). Burglary is chosen as an example for two reasons. First, with Henk Elffers and other members of an international collaboration, I am currently developing and testing models of burglar targeting strategies. One of the simpler models under development is used as an example here.

Second, for the crime of burglary the targets are fixed and hence the model is simpler than it would be for crimes such as robbery. For robbery, victimisation occurs when a victim and offender converge in space and time, in the absence of a capable guardian (Cohen and Felson, 1979). Thus, to model patterns of robbery, not only must the behaviour of offenders and guardians against crime be simulated, but we must also model the routine activities of the remainder of the population. To generate such a simulation that has an acceptable level of ecological validity will require considerable work. Thus, a simple model is used here to illustrate the points of central importance.

A picture tells a thousand words. Watching and playing with a simulation tells considerably more. Thus, rather than just describing the model developed and the results generated, the simulation was developed in the NetLogo programming language (Wilensky, 1999) discussed above to allow readers to download and play with it. The model developed for this section is also free to download.[10]

The simulation is made up of two basic elements; the world and two classes of agents (police and offenders) that move around it. The virtual world is made up of a grid of regular sized cells or, to use the nomenclature of NetLogo, patches which represent crime opportunities; homes in this model. Each home is assigned a crime attractiveness value to represent its risk of victimisation. A choice of different models are available that simulate different patterns in the variation in risk across homes. These range from a homogeneous surface where every home has the same risk; a binary risk surface where homes in the East and West side of the grid have different risks; a quad in which risk varies in the four corners of the surface; and, a surface over which the variation in risk is generated using a combination of a uniform RNG and a smoothing function. Figure 4 shows examples of two of the types of surface discussed. The user with a little programming knowledge can easily generate more, and real data on the spatial variation in crime risk could be imported.

INSERT FIGURE 4 ABOUT HERE

Agents are used to model the activity of offenders and police officers. Each agent can traverse the virtual world according to a set of predefined rules. For example, agents have vision for up to *n* cells ahead (this parameter is currently homogeneous within each agent set and selected by the user). Where

[10] http://www.jdi.ucl.ac.uk/british_academy_network/history/index.php.

an agent is at the edge of the grid, they see only as many cells ahead as exist. The number of agents, the number of time steps, how far the agents can see and many other variables can be varied. In the following sections a little more detail is provided about each class of agent.

Offender Agents

For every time step, each offender agent turns in the optimal direction, moves and then decides whether or not to commit an offence. The decision to commit a crime is a function of a number of factors which will be discussed below but also of the proximity of police agents. If a police agent is within five cells (this may of course be varied but is held constant here) of an offender agent, the offender agent will not commit a crime. This is used to model the preventive effect of police patrols.

Offender agent movement is determined in the following way:

1. Each agent looks directly ahead in all directions in its Von Neuman neighbourhood (North, East, South and West) and calculates a value to represent the cumulative opportunities in that direction. For example, when looking North, the agent will look at the *n* cells North of it, and add the attractiveness values for each cell together.

   This evaluation has two important additional features. First, every agent's vision is locally weighted so that greater importance is given to proximate locations. Second, rather than being a purely objective calculation, to generate a degree of bounded rationality (see Cornish and Clarke, 1985), a random number is added (or subtracted) to the cumulative value for each direction considered. The maximum value of the random number generated is a function of the risk ahead; the random number will tend to be largest when the attractiveness of the opportunities ahead is greatest.

2. The agent also considers how much crime has been committed in each direction with a view to avoiding areas that have been over-foraged (see Johnson and Bowers et al., 2008) and hence may be subject to police attention.

3. To determine the vector of travel, the agent turns in the direction that offers the optimal balance between the highest estimated cumulative opportunity value and the risk of encountering a police patrol at that time.

The agent then moves one cell and decides whether or not to commit a crime at that location. This is a two stage process, as follows:

a) The agent's state of readiness to offend (e.g., Clarke and Cornish,1985) at each moment is modelled as a function of an RNG and the individual offending rate for that offender (lambda). For the current model the value of lambda is held constant across agents and time, but this may be varied to examine the effect(s) of population heterogeneity or desistance decisions (Clarke and Cornish, 1985).

b) If the agent is in a state of readiness to offend, and there are no police agents nearby, the attractiveness of the cell in which the agent is located is evaluated. The higher the attractiveness of the particular cell, the higher the likelihood that the agent will offend at that location. Whether an offence occurs or not is thus determined by the attractiveness of the location, the output of an RNG and the agent's interaction with the police agents. When an offence takes place at a particular location, the number of crimes recorded at that cell is updated.

Thus, the offender agents' behaviour is a function of a series of simple rules. As discussed, the model here used is deliberately simple to illustrate the main points. However, the interesting thing is that the emergent behaviour need not be so simple and the local interactions between the different classes of agents and their environment can generate patterns that readers may expect would require more complex models.

Police Agents

Two types of patrolling strategy are modelled. For the first, police agents move randomly. In this case, for each time step, every police agent randomly selects a direction of travel from its Von Neuman neighbourhood and then (on the basis of the output of an RNG) moves one or two cells in that direction. The police agents could of course be limited to moving only one cell per time cycle, but they are not in this model.

In the second model, police agents decide where to patrol based on the volume of crime they "see" in front of them. This is essentially a hotspot policing model. Each police agent looks *m* cells ahead in each direction of its Von Neuman neighbourhood and evaluates the volume of crime in every direction. The value of *m* is selected by the user. Different values are likely to affect the performance of the model but here *m* is set to a constant value of 40 across simulations. The parameter *m* was set in this way so that the police agents would be able to "see" further than the offender agents. This was to reflect the fact that in the real world, the police may be advised about crime locations by other officers or by a command and control centre. Having evaluated the volume of crime in each direction, the police agent moves one or two cells in the direction in which most crime has occurred. The agents act independently and do not communicate with each other. Nor do they attempt to avoid directions where there are other police agents.

Initial Conditions

At the start of the simulation, each agent is placed at a random location on the grid. For the offenders this may thought of as their home location. Considering the variation in victimisation risk across homes, for the current model, the grid is divided into four quadrants across which (but not within) the attractiveness of targets vary.

Results of the Policing Simulation

The results of a simulation may be analysed in much the same way as data generated in the real world. For example, Figure 5 shows an example of a hotspot map generated by one run of a simulation.

INSERT FIGURE 5 ABOUT HERE

To provide a more systematic analysis, for each model described, the effect of increasing the number of police agents is tested on the dependent variable – the volume of offences committed. Specifically, models are tested with 0, 10, 20, 40 and 80 police agents. Every model is executed 20 times and for 1,000 time steps. The reason for running the model a number of times is that the initial conditions of the model (where the agents are initially located) will influence the outcome of the simulation.

Figure 6 shows the mean number of crimes committed for each model tested. With the exception of the model for which there are zero police agents ($p=0.79$), less crime is committed for the hotspot policing model (all $ps<0.01$). For the hotspot policing model, increasing the number of police officers has a roughly linear effect on the volume of crime committed. For the random patrol model, the effect is more variable.

INSERT FIGURE 6 ABOUT HERE

Assumptions and Parameter Sweeps

It should be obvious from the above discussion that many decisions regarding parameter settings and the formalisation of decision rules are required to produce a working simulation. On the one hand, this is useful as it forces the researcher to explicitly formalise rules and consider issues that they otherwise might not. On the other, this reminds us of one of the potential problems with simulation. When so many parameters can be manipulated and processes implemented in a range of ways, each permutation will require evaluation if conclusion validity is to be established in any meaningful way.

Fortunately, modelling platforms such as Netlogo have tools which enable the testing of different permutations of parameter values. Running each configuration of parameter settings also enables the effects of chance and initial conditions to be examined on the output of a model.

For the above results, the findings are of course only of any utility if the underlying model reflects the way the world is. There is no suggestion here that the model does. However, whilst considerable caution is clearly necessary regarding the ecological validity of simulation models, the promise is equally exciting. Even at this stage of maturity, simulation can be used for theory testing and falsification and, in the context of policy simulation, could (with some degree of caution) be used to identify those conditions under which certain types of intervention –with simple crime reduction mechanisms – might work. Thus, simulation models could be used as tools to help researchers and practitioners think through how an intervention might work in a systematic way.

What Else Might be Included in the Simulation Model?

The simulation described above was quite simple. Further models could easily be developed to generate a more realistic simulation. The possibilities are considerable, but some examples include:

1) The incorporation of a street network which would affect agent mobility in different directions (Birks et al., 2008; Groff, 2007ab);
2) The inclusion of an algorithm to allow the agents to navigate the world in a more deliberate way (see, Groff, 2008). This could draw on findings from research concerned with space syntax (e.g., Hiller, 2004) which suggest that the geometry of the street network can affect navigation in subtle ways. For example, as a consequence of variation in sight lines due to street network configuration, the route taken from two locations A and B, is often different from the route taken from B to A by the same person;
3) Barriers to offender movement, physical or social (e.g., Bernasco and Nieuwbeerta, 2005), could be included;

4) Offenders could form individual mental maps of their world based on their learning experiences in it, and perhaps choose to offend with a higher probability in those areas they know best;
5) A range of routine activity nodes (e.g., home, friends homes, shops) could be generated for each agent (e.g., Groff, 2007a) and they could be encouraged to travel to them on certain days and at particular times. More detailed insights into human mobility patterns, routine activities and the anisotropy of human movement could also be incorporated into the models (Gonzalez et al., 2008);
6) A form of gravity could be used to encourage offender agents to avoid travelling too far from their home location or routine activity nodes;
7) The offender agent's readiness to offend could be varied over time, perhaps as a function of recent success; and,
8) The activity and visibility of the police could be modelled in a more realistic way, drawing upon research extant (e.g., Clarke and Hough, 1984).

Other Types of Intervention?

This type of approach could, of course, be used to look at other types of crime prevention strategies, such as street lighting. However, in some cases there may be a risk of producing models that generate nothing more than pre-supposed emergence. To illustrate, consider that if we produce a model in which we say that street lighting reduces the risk of victimization by 20% and then run the model and examine the impact of intervention, the problem will be that the model is tautological. The results of the simulation will only tell us what we put into it; it will not test a causal mechanism.

As an alternative, we could test the effects of increased illumination on crime reduction mechanisms. For example, we could see what the effect of increasing an offender's visible range has on their foraging strategies and rate of offending given the possible changes in attendant risks. Likewise, the effects of increased illumination on the potential victims' visible range, and hence potential avoidance strategies, could be tested. However, it is worth noting that such a mechanism assumes that the effect of street lighting is produced as a consequence of the increase in luminescence generated. In relation to this, an illuminating finding is provided by Farrington and Welsh (2002, p. 36) who examined the effects of street lighting on crime during the day and night. Similar effects were observed at both times of the day, suggesting that the mechanism of change was not simply the change in illumination.

CONCLUSION

The potential benefits of using computational methods for theory testing and evaluation are numerous. The examples provided in this chapter were intended to illustrate some of the possibilities and how they might be used to inform the sequenced decision making processes that policy actors routinely engage in. It is important to realize that other types of simulation methodology exist and that the types of questions considered here represent the tip of the iceberg. However, it is also important to remember that care needs to be taken when using simulations. The results of a model do not tell us what will happen in the real world, only what might happen if the assumptions of the model are valid. They may not be, but numbers can be seductive and so one risk is that the output of simulation models may inspire confidence where caution would be more appropriate. Before policy simulation is pursued in earnest, the way forward requires considerable effort in testing and maximizing the validity of models designed to simulate normal patterns of crime, perhaps along the lines suggested by Townsley and Johnson (2008). For individual studies, it will be important to conduct sensitivity analyses by sweeping the parameter space of models used and summarizing the effect of changes made on the outcomes generated. Moreover, replication is the cornerstone of good science and so the independent verification of findings from research which involves simulation will be important (Townsley and Johnson, 2008).

Despite the obvious need for caution, as other authors have concluded (e.g., Groff and Birks, 2008) simulation methods offer researchers exciting new tools for research. If appropriately specified, simulations could allow the systematic testing of ideas before fieldwork begins, as well as providing tools to help decision makers better understand the likely impacts of tested interventions.

♦

Address for Correspondence: UCL Jill Dando Institute of Crime Science, University College London , Second Floor Brook House, 2 - 16 Torrington Place, London WC1E 7HN, UK; e-mail: shane.johnson@ucl.ac.uk

Acknowledgments: This work was supported by a British Academy research grant, award number LRG-45507. I would like to thank the editors and the other contributors to this book each of whom provided helpful comments on an earlier draft of this paper. In particular I would like to thank Gary Henry, Paul Ekblom and Elizabeth Groff for detailed comments and suggestions. Thanks also go to Ken Pease and Henk Elffers.

REFERENCES

Agresti, A. (2002). *Categorical Data Analysis (2nd Edition)*. Florida: Wiley.

Alimadad, A., Borwein, P., Brantingham, P., Brantingham, P., Dabbaghian-Abdoly, V., Ferguson, R., Fowler, E., Ghaseminejad, C.G., Li, J., Pollard, N., Rutherford, A. & Waall, A. (2008). Using Varieties of Simulation Modelling for Criminal Justice System Analysis. In L. Liu & J. Eck (Eds.), *Artificial Crime Analysis Systems: Using Computer Simulations and Geographic Information Systems*. New York: Information Science Reference.

Bak, P. (1996). *How Nature Works: The Science of Self Organized Criticality*. New York: Springer-Verlag.

Barr, R. & Pease, K. (1992). A Place for Every Crime and Every Crime in its Place: An Alternative Perspective on Crime Displacement. In D.J. Evans, N.R. Fyfe & D.J. Herbert (Eds.), *Crime Policing and Place: Essays in Environmental Criminology*. London: Routledge.

Batty, M. (2007). *Cities and Complexity*. Cambridge, MA: MIT Press.

Bernasco W. (2007). Them again? Same-offender involvement in repeat and near repeat burglaries. *European Journal of Criminol*ogy,5, 411-431..

Bernasco, W. & Nieuwbeerta, P. (2005). How Do Residential Burglars Select Target Areas? *British Journal of Criminology*, 45, 295-315.

Birks., D.J., Donkin, S. & Wellsmith, M. (2008). Synthesis over Analysis: Towards an Ontology for Volume Crime Simulation. In L. Liu & J. Eck (Eds.), *Artificial*

Crime Analysis Systems: Using Computer Simulations and Geographic Information Systems. New York: Information Science Reference.

Bowers, K. J. & Johnson, S. D. (2005). Domestic burglary repeats and space-time clusters: The dimensions of risk. *European Journal Criminology,* 2, 67-92.

Bowers, K.J., Lab, S. & Johnson, S.D. (2008). Evaluating Crime Prevention Using Survival Analysis. *Policing: A Journal of Policy and Practice*, 2, 218-225.

Bowers, K.J., Johnson, S.D., & Hirschfield, A.F.G. (2005). Closing off Opportunities for Crime: An Evaluation of Alley-gating. *European Journal on Criminal Policy and Research*, 10(4), 285-308.

Bowers, K.J. & Johnson, S.D. (2005). Domestic burglary repeats and space-time clusters: the dimensions of risk. *European Journal of Criminology*, 2(1), 67-92.

Bowers, K.J., Johnson, S.D. & Pease, K. (2005). Victimisation risk, housing type and area: the ecological fallacy lives! *Crime Prevention and Community Safety: An International Journal,*7(1), 7-18.

Bowers, K.J. & Johnson, S.D. & Hirschfield, A.F.G. (2004). The measurement of crime prevention intensity and its impact on levels of crime. The *British Journal of Criminology*, 44(3), 1-22.

Braga, A.A. (2001). The Effects of Hot Spots Policing on Crime. *Annals of the American Academy of Political and Social Sciences*, 578, 104-125.

Braga, A.A. & Bond, B.J. (2008). Policing Crime and Disorder Hot Spots: A Randomized Controlled Trial. *Criminology,* 46(3), 577-607.

Brantingham, P. & Brantingham, P. (1993). Environment, routine, and situation: Toward a pattern theory of crime. In R.V. Clarke & M. Felson (Eds.), *Routine Activity and Rational Choice*,. New Brunswick, NJ: Transaction Publishers.

Budd, T. (1999). Burglary of domestic dwellings: Findings from the British Crime Survey. *Home Office Statistical Bulletin*, 4/99. London: Home Office.

Campbell, D.T. (1957) Factors relevant to the validity of experiments in social settings, *Psychological Bulletin*, 54:297-312

Campbell, D.T. & Stanley, J.C. (1963). *Experimental and Quasi-Experimental Designs for Research*. Boston: Houghton Mifflin.

Clarke, R.V. & Cornish, D.B. (1985). Modelling Offenders' Decisions: A Framework for Research and Policy. In M. Tonry & N. Morris (Eds.), *Crime and Justice: An Annual Review of Research*, 6. Chicago: University of Chicago Press.

Clarke, R.V. & Hough, M. (1984). Crime and Police Effectiveness. *Home Office Research Study 79*. London: Home Office.

Cornish, D. & Clarke, R. V. (2003). Opportunities, Precipitators and Criminal Decisions: A Reply to Wortley's Critique of Situational Crime Prevention. In M. J. Smith & D.B. Cornish (Eds.), *Crime Prevention Studies*, vol. 16,. Monsey: Criminal Justice Press.

Dray, A., Mazerolle, L., Perez, P. & Ritter, A. (2008). Drug Law Enforcement in an Agent-Based Model: Simulating the Disruption. In L. Liu & J. Eck (Eds.), *Artificial Crime Analysis Systems: Using Computer Simulations and Geographic Information Systems*. New York: Information Science Reference.

Eck, J. (1993). The Threat of Crime Displacement. *Criminal Justice Abstracts*, 25, 527-546

Ekblom, P. (2000). The Conjunction of Criminal Opportunity – a tool for clear "joined-up" thinking about Community Safety and Crime Reduction. In S. Ballintyre, K. Pease & V. McLaren (Eds.). *Secure Foundations: Key Issues in Crime Prevention, Crime Prevention and Community Safety*. London: Institute for Public Policy Research.

Ekblom, P., Law, H, & Sutton, M. (1996). Safer Cities and Domestic Burglary. *Home Office Study 164*. London: Home Office.

Ekblom, P. & Pease, K. (1995). Evaluating Crime Prevention. In M. Tonry & D Farrington (Eds.), *Building a Safer Society: Strategic Approaches to Crime Prevention, Crime and Justice: A Review of Research*, 19. Chicago: University of Chicago Press.

Farrell, G. (2005). Progress and prospects in the prevention of repeat victimization. In N. Tilley (Ed.), *Handbook of Crime Prevention and Community Safety*. Cullompton, UK: Willan.

Farrington, D.P. & Welsh, B.C. (2006). How Important is "Regression to the Mean" in Area-Based Crime Prevention Research. *Crime Prevention and Community Safety*, 8, 50-60.

Farrington, D.P. & Welsh, D.C. (2002). *Effects of Improved Street Lighting: A Systematic Review*. Home Office Research Study 215. London: Home Office.

Farrington, D.P., Coid, J.W., Harnett, L., Jolliffe, D., Soteriuo, N., Turner, R. & West, D.J. (2006). Criminal Careers up to age 50 and life success up to age 48: new findings from the Cambridge Study in Delinquent Development. *Home Office Research Study 299*. London: Home Office.

Felson, M. (2005). *Crime and Nature*.Thousand Oaks: Sage.

Forrester, D., Chatterton, M. & Pease, K. (1988). *The Kirkolt Burglary Prevention Project, Rochdale*. Crime Prevention Unit Paper 13. London: Home Office.

Gilbert, N. & Troitzsch, K.G. (2005) *Simulation for the Social Scientist*, 2nd ed.. Maidenhead, UK: Open University Press

Gonzalez, M.C., Hidalgo, C.A. & Barabasi, A-L. (2008). Understanding Individual Human Mobility Patterns. *Nature*, 453, 779-782.

Groff, E. R. (2007a). Simulation for Theory Testing and Experimentation: An Example Using Routine Activity Theory and Street Robbery. *Journal of Quantitative Criminology,* 23(2), 75-103.

Groff, E. R. (2007b). 'Situating' Simulation to Model Human Spatio-Temporal Interactions: An Example Using Crime Events. *Transactions in GIS,* 11(4), 507-530.

Groff, E. R. (2008). Spatio-Temporal Aspects of Routine Activities and the Distribution of Street Robbery. In L. Liu & J. Eck (Eds.), *Artificial Crime Analysis Systems: Using Computer Simulations and Geographic Information Systems*, pp. 226-251. Hershey, PA: Idea Group.

Groff, E. & Birks, D. (2008). Simulating Crime Prevention Strategies: A Look at the Possibilities. *Policing: A Journal of Policy and Practice*. 2, 175-184.

Groff, E. R. & Mazzerole, L. (2008). Simulated Experiments and their Potential Role in Criminology and Criminal Justice. *Journal of Experimental Criminology, 4*(3). I WAS REFERENCING THE WHOLE VOLUME HERE!!

Hillier, B. (2004). Can Streets be Made Safe? *Urban Design International*, 9, 31-45.

Johnson, S.D. (2008). Repeat Burglary Victimization: A Tale of Two Theories. *Journal of Experimental Criminology*,4, 215-240.

Johnson, S.D. & Bowers, K.J. (2004a). The burglary as clue to the future: the beginnings of prospective hot-Spotting. *European Journal of Criminology*, 1(2), 237-255.

Johnson, S.D. & Bowers, K.J. (2004b). The stability of space-time clusters of burglary. *British Journal of Criminology*, 44(1), 55-65.

Johnson, S.D., Summers, L. & Pease, K. (2008). Offender as Forager: A Direct Test of the Boost Account of Victimization. *Journal of Quantitative Criminology*, online first DOI 10.1007/s10940-008-9060-8

Johnson, S.D., Bowers, K.J., Birks, D. & Pease, K. (2008). Predictive Mapping of Crime by ProMap: Accuracy, Units of Analysis and the Environmental Backcloth, D. Weisburd, W. Bernasco & G. Bruinsma (Eds.), *Putting Crime in its Place: Units of Analysis in Spatial Crime Research*. New York: Springer.

Johnson, S.D., Bernasco, W., Bowers, K.J., Elffers, H., Ratcliffe, J., Rengert, G. & Townsley, M.T. (2007). Near Repeats: A Cross National Assessment of Residential Burglary. *Journal of Quantitative Criminology*, 23(3), 201-219.

Johnson, S.D., Bowers, K.J., Jordan, P., Mallender, J., Davidson, N. & Hirschfield, A.F.G. (2004). Estimating crime reduction outcomes: How many crimes were prevented? Evaluation: the International Journal of Research and Practice, 10, 327-348.

Lacoste, J. & Tremblay, P. (2003). Crime and Innovation: A Script Analysis of Patterns in Check Forgery. In M. J. Smith & D.B. Cornish (Eds.), *Crime Prevention Studies*, vol 16. Monsey: Criminal Justice Press.

Liang, J., Liu, L. & Eck, J. (2001). *Simulating Crimes and Crime Patterns Using Cellular Automata and GIS*. Presentation to the 2001 UCGIS Summer Assembly. Buffalo, New York.

Liu, L. & Eck, J. (2008). *Artificial Crime Analysis Systems*. New York: Information Science Reference.

Marchant, P. (2005). What Works? A Critical Note on the Evaluation of Crime Reduction Initiatives. *Crime Prevention and Community Safety*, 7, 7-13.

North, B.V., Curtis, D. & Sham, P.C. (2002). A note on the calculation of empirical p values from Monte Carlo procedures. *American Journal of Human Genetics,* 71, 439-441.

Openshaw, S. (1984). *The Modifiable Areal Unit Problem*. Norwich: Geo books.

Pawson, R. & Tilley, N. (1997). *Realistic Evaluation*. London: Sage.

Pease, K. (1998). Repeat victimization: taking stock. Home Office Police Research Group, Crime Detection and Prevention Series, Paper 90. London: Home Office.

Polvi, N., Looman, T., Humphries, C. & Pease, K. (1991). The time course of repeat burglary victimization. *British Journal of Criminology,* 31**,** 411-414.

Pyke, G.H. (1981). Honeyeater Foraging: A Test of Optimal Foraging Theory. *Animal Behaviour*, 29, 878-888.

Ratcliffe, J.H. & Rengert, G.F. (2008) Near repeat patterns in Philadelphia shootings, *Security Journal*, 21, 58-76.

Shadish, W.R., Cook, D. & Campbell, D.T. (2002). *Experimental and Quasi-Experimental Designs for Generalized Causal Inference*. Boston: Houghton Mifflin.

Sherman, L. (2007). The Power Few: Experimental Criminology and the Reduction of Harm. *Journal of Experimental Criminology*, 3, 299-321.

Summers,L. Johnson, S.D. & Pease, K. (2007). El Robo de (Objetos en) Vehículos y su Contagio a través del Espacio y el Tiempo: Aplicaciones de técnicas epidemiológicas. *Revista Electronica de Investigacion Criminologica*, 5, 1-22.

Tabachnick, B.G. & Fidell, L.S. (2001). *Using Multivariate Statistics*. Boston: Allyn and Bacon.

Tilley, N. (1993). After Kirkholt- Theory, Method and Results of Replication Evaluations. *Police Research Group Crime Prevention Unit Series Paper 47*. London: Home Office.

Townsley, M. & Johnson, S.D. (2008). The Need for Systematic Replication and Tests of Validity in Simulation. In L. Liu & J. Eck (Eds.), *Artificial Crime Analysis Systems: Using Computer Simulations and Geographic Information Systems*. New York: Information Science Reference.

Townsley, M., Homel, R. & Chaseling, J. (2003). Infectious burglaries: A test of the near repeat hypothesis. *British Journal of Criminology,* 43, 615-633.

Visher, C. & Linster, R. (1990) A Survival Model of Pretrial Failure. *Journal of Quantitative Criminology* 6, 153-184.

Willensky, U. (1999). *Netlogo* (http://ccl.northwestern.edu/netlogo). Center for Connected Learning and Computer-Based Modelling, Northwestern University. Evanston, IL.

Weisburd, D., Bushway, S., Lum, C. & Yang, S-M. (2004). The Trajectories or Crime at Places: A Longitudinal Study of Street Segments in the City of Seattle. *Criminology*, 2, 283-321.

Table 1. Estimates of the Volume of Crime for Three Different Models (Generated from 100 Iterations of the Simulation)

	Mean	Variance	SD
Area-level model (1)	218.9	229.6	15.2
RV model (2)	218.2	279.0	16.7
Space-Time model (3)	218.3	371.9	19.3

(1) Mixed Poisson model.
(2) As for (1) but the risk to victimised homes is temporarily elevated.
(3) As for (2) but the risk to nearby homes is also temporarily elevated.

Table 2. Simulated Effects of Using Different Dosages of Implementation Assuming a 50% Crime Reductive Effect at the Individual Level and Only Chance Variation in Risk Over Time (Generated from 100 Iterations of the Simulation)

	Mean % Reduction	Min % Reduction	Max % Reduction
Abrupt Implementation			
100% of homes	49.0	38.5	63.3
50% of homes - Low Risk	14.6	-3.6	28.9
50% of homes - High Risk	34.5	22.0	47.7
Gradual Implementation			
100% of homes	30.6	16.9	44.0

Figure 1. Survival Times for Target Hardened and Non-Target Hardened Homes (Adapted from Bowers et al., 2008)

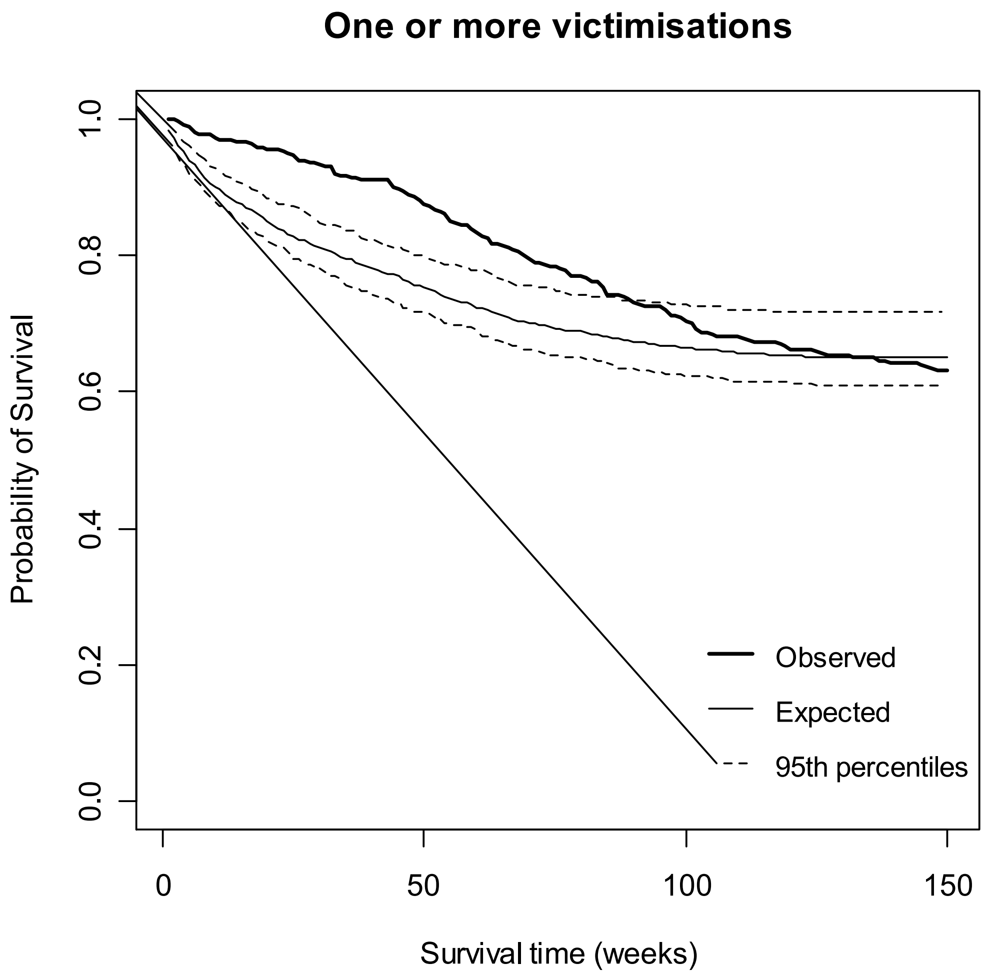

Figure 2. Area Selected for Microanalytical-Simulation Experiment (Ordnance Survey © Crown Copyright. All Rights reserved)

Figure 3. Examples of Kernel Density Maps to Show the Patterns Generated for a Fragment of the Virtual World for Two Realisations of the Simulation Experiment (Ordnance Survey © Crown Copyright. All Rights reserved)

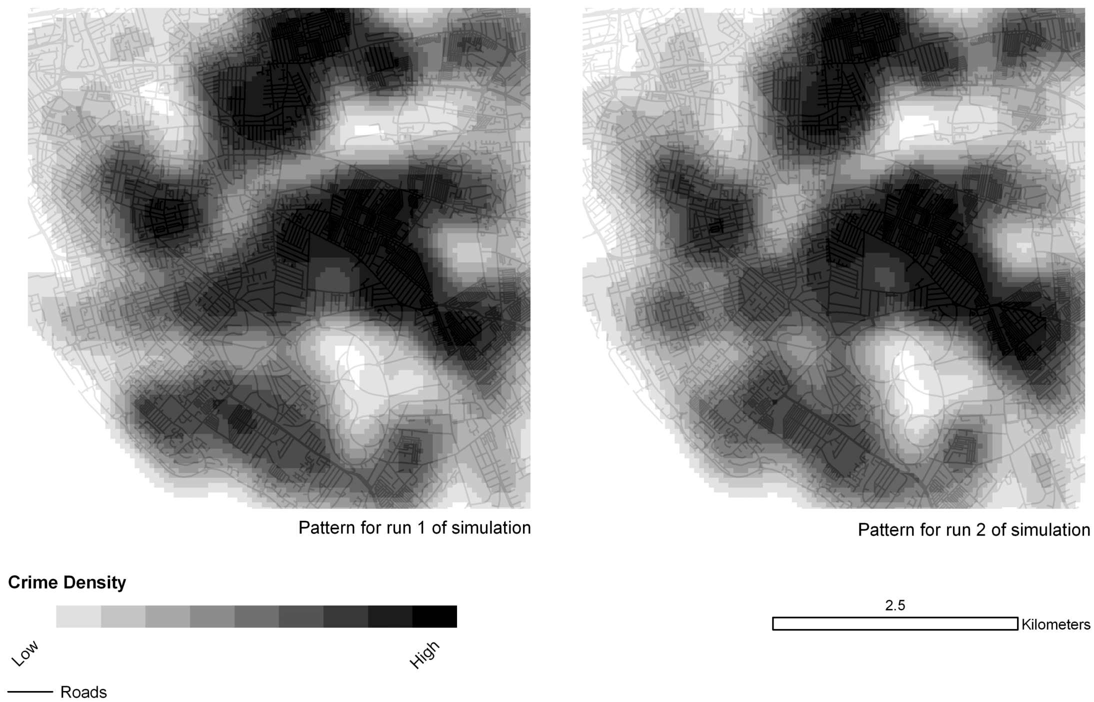

Figure 4. Example Risk Surfaces (Left: binary surface, Right: Random variation)

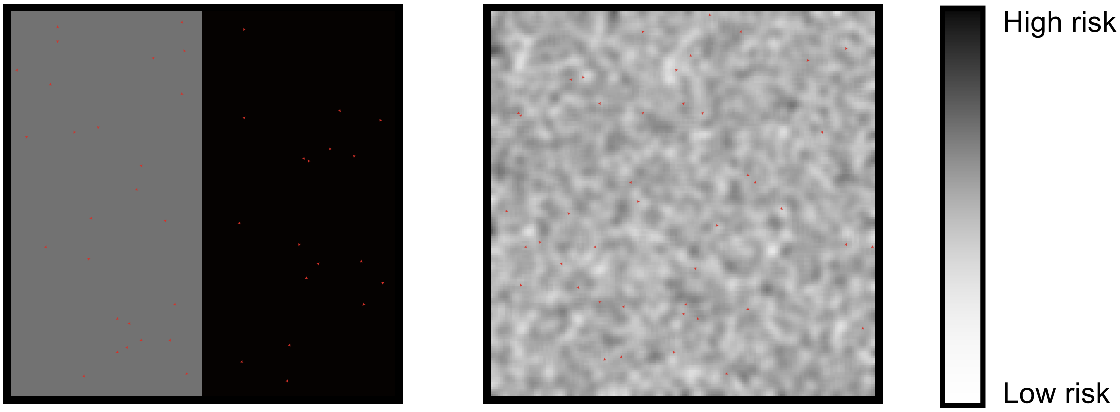

Figure 5. An Example of a Kernel Density Map Generated by an Agent Based Simulation.

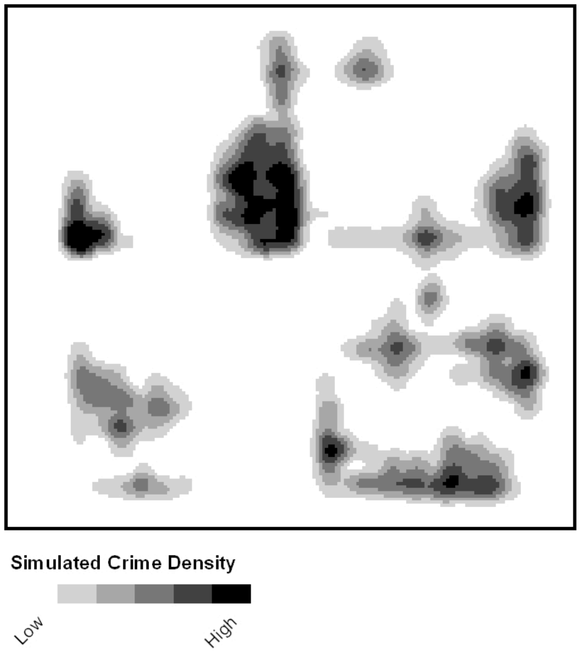

Figure 6. Impact on Crime of two Patrolling Strategies and Changes in the Number of Police Agents (20 simulations per model, 1000 time cycles per simulation).

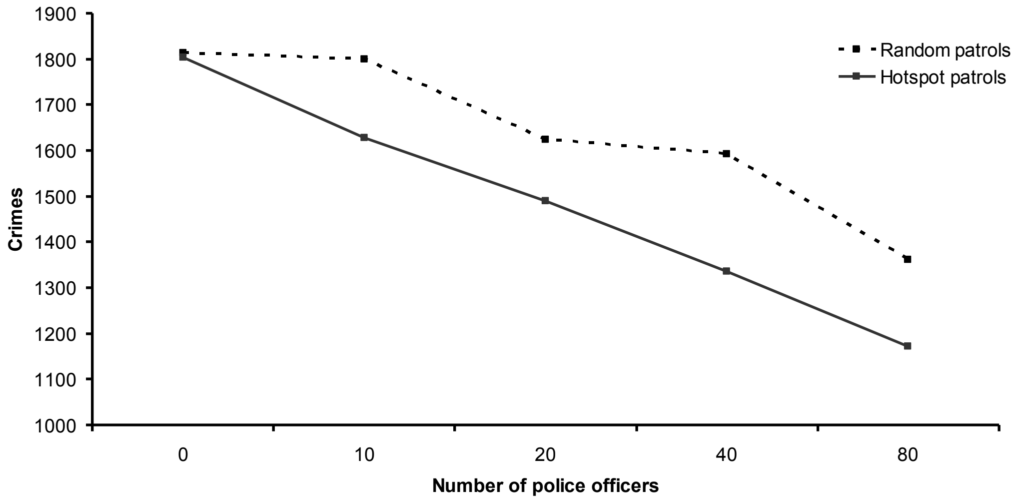

Box 1 – Microanalytical Simulation

Microanalytical simulation is used to model how the characteristics of a population might change over time under different conditions. Gilbert and Troitzsch (2003) provide a number of examples of how this approach may be used to answer different types of policy questions. One example considers how a researcher might estimate the likely future demands on the nursing sector by computing the likely population of those aged 60 or more with and without support networks (e.g. close relatives). For example, for a given (real) population, this type of model can be used to simulate how the composition of the overall population, and people's support networks, will change when one takes account of factors such as likely variation in rates of births, deaths, divorce, the effects of chance, and so on.

For this type of model, patterns can be measured at the level of the individual or at the aggregate level of the population. This approach requires detailed data on the population considered as well as the specification of those processes that may affect the population concerned.

Box 2 - Agent Based Simulation

Agent based models (sometimes also referred to as Individual based models) have been used to study a range of phenomenon including the spread of epidemics, consumer behavior, and pedestrian movement. For social science applications, agents, which may be thought of simple representations of people (or whatever is being simulated), are programmed with simple rules about how to behave in a simulated environment. When a simulation is running, agents make decisions based on the rules specified but their choices are not pre-determined, and are instead influenced by the decisions of other agents as well as the effects of chance. Thus, even though the rules provided may be simple, when agents interact with each other or adapt to the impacts that other agents have had on the environment, the behaviour that *emerges* can be complex and unanticipated by the researcher. Such complexity cannot easily be modeled (if at all) using other methods.

Those interested in playing with a range of agent based models may do so by downloading the NetLogo program, available at: http://ccl.northwestern.edu/netlogo/.

For those interested in reading about how such models can be used in a policy context, Batty (2007) provides a fascinating example of how an Agent based simulation was used to examine the likely influence of changes in pedestrian routes on levels of over-crowding at the Notting Hill carnival in London.