King's Fund

MODELLING URGENT EVENTS

Seán Boyle
Anthony Harrison
Jeff Tansley
Jane Bryan-Jones

January 1998

Abstract

A major policy concern for the NHS is how to deal with the increase in emergency hospital admissions widely observed in the 1990s. There is a burgeoning literature on the phenomenon and factors associated with it (Kendrick, 1996; Edwards, 1997). The impact on all aspects of health care provision from longer waiting times for elective care, to more stress on primary care services and staff generally, to poorer outcomes for people, is also recognised. This paper presents a systems approach to understanding the way in which unpredictable events are dealt with by health care services. It does this by developing a model of 'urgent' events and simulating the workings of an actual system of care.

The purpose of this work is threefold: to enable and encourage mutual learning about a complex system, both at an institutional level - the hospital for example - and a general system-wide level; to indicate where knowledge is lacking and hence show clear requirements for further detailed research; and, finally, to allow particular policy options to be tested within a well-defined assumptive framework before actual implementation takes place. The paper is in three parts. The first outlines the issues addressed, setting these in the context of the crisis in NHS emergency provision from which the impetus for this work arose. The second provides a general discussion of models of perceived need and response which lie at the heart of the simulation process.

The final section of the paper moves on from this general discussion to the simulation model itself, providing a description of the simulation objects. It is in the nature of this type of work that it is constantly developing as different phases of the model are adjusted, calibrated and improved. This paper therefore reflects an interim report on the work done to date rather than its ultimate conclusion.

Seán Boyle & Anthony Harrison

King's Fund
11-13 Cavendish Square
London
W1M 0AN
TEL 0171 307 2400
FAX 0171 307 2440
E-MAIL inner@kehf.org.uk

Jeff Tansley & Jane Bryan-Jones

An Teallach Limited
Laggan Bridge
Newtonmore
PH20 1BT
TEL 01528 544 312
FAX 01528 544 360
E-MAIL jeff@an-teallach.com

CONTENTS

1 Rising Demand for Emergency Care

The prompt provision of emergency care is perhaps the most important requirement of any health care system, at least in the mind of the users who wish to know that the system will respond effectively if their lives are in danger or they are faced with a less serious threat that nevertheless requires a rapid response. In recent years NHS provision of emergency care has come under frequent criticism for failing to meet even basic standards relating to availability and time waiting in inappropriate circumstances.

There is a considerable body of evidence showing the NHS faces unmanageable surges in emergency admissions particularly in the winter months, rises in A&E attendances and increased demand on ambulance services leading to delays and cancellations in planned health care work. These reports range from detailed research studies to press coverage of individual cases.

Between 1990/91 and 1994/95 there was an increase of 20 per cent - 4.7 per cent per annum - in the level of emergency admissions to hospitals in England. In itself this is not an inordinate increase but it occurred at a time when the slack within the hospital system has steadily declined as the drive to improve efficiency, mainly through increased day case rates, has meant that the system has often appeared overloaded.

Moreover, the average increase masks considerable variation, both between individual hospitals and areas, and within the year. Each winter there is a crisis in many hospitals throughout the UK. London is particularly badly affected. These crises manifest in long waits to be seen in accident and emergency departments, overnight stays on trolleys because of bed shortages, and even closures of hospitals to emergency admissions.

At the same time there has been an even greater increase in the use of ambulance services. Between 1992/93 and 1994/95 demand for ambulance services rose by over seven per cent per annum, both in London and in England as a whole. More anecdotally, health care professionals from junior doctors to GPs and district nurses report working under greater pressure though this is difficult to quantify or tie to particular factors.

There appears however to be no simple explanation for these increases. The indications are that a range of factors are at work including an ageing population, environmental and socio-economic conditions, rising expectations of what the health service should provide, and changes in GP provision of out-of-hours

services.

These pressures along with other developments, particularly in information technology, have led to a large amount of experimentation with new ways of providing emergency care. These include minor injuries units - which may be linked by IT to large hospitals - new ways of organising and staffing A&E departments and new ways of providing ambulance services.

Although the immediate pressures on the NHS may be most acute on the hospital A&E department, the key to managing the pressures better must lie in looking at the role of all providers of emergency care, the context in which they work and the way users access the services available to them.

This paper therefore develops an approach which recognises that the provision of emergency care by hospitals is just one part of an extremely complex system of response to urgent need for health care. To do this it attempts to map and simulate the interactions of a range of individuals and provider agents within a single emergency care system. Only in this way will it be possible to understand the knock-on effects of policy changes in one provider on other parts of the system.

It is of course possible to consider elements of the overall system in isolation. Indeed, this is often necessary in gaining some initial understanding. However, we believe that the problems referred to above arise from the fundamental complexity of the system of health care provision, that they are best addressed through an overall understanding of the way in which individuals and provider agencies interact, and that the development of national policies to deal with perceived problems in one sector - hospital emergency provision for example - are stymied by this failure of understanding.

This view of the nature of the problem leads to the solution described in this paper - the construction of a simulation model - which attempts to build a framework of understanding within which limited existing knowledge can be utilised. In this way we hope to enable and encourage mutual learning about a complex system, both at an individual provider and a general system-wide level; to indicate where knowledge is lacking and hence show clear requirements for further detailed research; and, finally, to allow particular policy options to be tested within a well-defined set of assumptions before actual implementation takes place.

In the next section we discuss the way individuals perceive their need for care

and response of those who become users to the way care is provided. It is a fundamental assumption of what follows that the demand for urgent care, both as a whole and for individual providers depends on the system of care itself.

This assumption rests on two considerations:

- individuals may be able to meet their own demand for care and not call on the formal, professional care system;
- the perceived need for care may disappear with time: knowing this from experience people may defer accessing the care system until they are more certain they require assistance.

In other words, how people respond to a perceived need for urgent care will depend on how they perceive the risks they face, the costs, including time and inconvenience, of seeking care and the information available to them about the options available.

This leads to a key distinction between the formal and the informal care system. The formal system comprises services such as hospitals, GPs, ambulances which provide or support the provision of urgent care. The informal system consists of individuals who perceive a need for care but provide it for themselves, and their family or friends who may also provide care as well as advice or simply reassurance.

If someone concludes they need professional help, then they may access the formal system by phone or by going to a health care facility of some kind. What they do and where they go depends on the nature of the formal care system - how far the nearest hospital is, what other options are available and so on.

The final section of the paper moves from general discussion to the simulation model itself. It is in the nature of this type of work that it is constantly developing as different phases of the model are adjusted, calibrated and improved. The paper should be seen as reflecting work in progress rather than a final product.

2 Modelling Perceived Need and Response

The model can be seen as comprising three main components. First, a description is required of the context which gives rise to demands on the formal care system in terms of numbers of individual cases, their different perceived and actual morbidity levels, and the timing of those demands. The approach used here is population-based, ie the model assumes a given population whose perceived needs give rise to demands on the formal care system. Because distance to the formal care system can be important, it also assumes a specific geographical arrangement of the population. We term this *the environment.*

Second, a description is required of the professional or formal care system which the individual encounters. There are two parts to this: the first consists of the organisation of care; the second comprises the relationships between parts of the organisation in both clinical and managerial terms, and the resources - physical and human, and also financial- which are available to it. This description may closely represent what is available in a particular area or it may be 'idealised' if the aim is to test out a system of care that does not actually exist anywhere. We term this *the formal system.*

Finally there are the results produced by the system which at its most simple can be some measure of system effectiveness such as lack of pressure on hospital beds or lack of delays in admissions. Similarly for doctor hours, theatre time or radiology services. More complexity is introduced by the recognition that the system has a primary aim of providing effective health care - whether at one extreme, the formal system is effective in savings lives to, at the other, whether users are satisfied with the response in terms of promptness etc, when they access the system.

Of course situations will always arise in which someone dies in circumstances where a different configuration of services and time may have prevented that. Ideally, the evaluation of an efficient system would always involve consideration of the trade-off between acceptable risk, level of service and cost.

2.1 The Environment

To understand better the nature of the factors which determine the demand for various forms of emergency care we use the notion of an 'urgent event'. This is defined as a situation where an individual may consider whether or not to seek urgent medical attention. We take this to include situations where the individual does not actively 'consider' seeking urgent medical attention but would benefit

from doing so. The event may be anything from waking up with a headache to collapsing with a massive coronary.

Not all urgent events warrant the attention of the formal care system. By using this wide definition of urgent events we allow for the potentially huge latent demand for emergency and other easy-access services which might be aroused quite unintentionally by a new service ostensibly introduced to relieve some of the pressure on hospitals - for example better access to advice.

The events are assumed to arise within a given geographical area which is divided up into parts termed zones. The characteristics of the population and its activities, eg travel choices, and of zones, eg deprivation level or industrial type, are assumed to determine the demand for urgent care arising from that area. As a result this will vary from zone to zone.

At its simplest the model requires only one location or site where formal care can be provided. If we knew that on average, for every 1,000 people, 200 demands for urgent care within the formal care system would arise a year, it would be a simple matter to move from census data to an estimate of overall demand at that site within the year.

However, if the level of demand is a function of the characteristics of the population, then we may need information on age, socio-economic group, household size etc. It might also depend on zone rather than household characteristics, eg main roads, factories. Thus statistical relationships might be required, as they are in transport models, between zone characteristics and the number of incidents requiring urgent care.

In practice, there will be at least two forms of provision available - hospital and GP surgery. Once there is more than one site where treatment might be offered, other considerations come into play: personal preferences, the characteristics of the services available at each site, and the nature of the transport/travel system joining zones with treatment sites. If there are two treatment sites, H and P, then which people go to may be determined by how easy it is to reach them, how quickly people expect to be treated, how accurately they determine their own need for treatment (if only H provides advanced treatment) and so on.

Although this is a simple framework, it requires a way of allowing for the impact of distance, travel cost and other determinants of demand for services at particular sites and indeed for formal care overall. While initially it may make sense to use simple rules - people go to the nearest site - the capacity to use more

complicated ones should be allowed for.

The framework set out earlier envisages that demand for formal care will vary according to the choices people make about whether to seek treatment or not. These choices will in part depend on the nature of the options available: how far away they are, whether telephone advice is available and so on. Such choices will concern not only whether to seek help, but also when to do so. This means that both the total demand, and/or its timing may be a function of the formal system of provision.

For example if a third treatment site is opened, P*, this will:

- divert demand from H and P;
- generate new demand, ie demand not previously entering the formal system because it reduces the deterrent effect of distance and other factors for some users;
- (possibly) shift the time profile of demand if P* is open at different times to the existing facilities.

So far, we have talked mainly of demand in general but of course the characteristics of that demand will vary greatly from the trivial to the most life threatening. If H and P, as well as P*, offer different facilities, then any one initial incident may give rise to demand at more than one site, eg if someone in need for care goes to P when only H can treat them. This matters in terms both of workload but also, in extreme cases, care outcome. The requirement here is for onward referral rules and a 'counting' mechanism which allows for multiple stages in one episode of use.

Although the time at which an urgent event leads to a demand on the formal care system may be influenced by personal factors and by the care system itself, in general the demand on the system cannot be postponed indefinitely as for example the demand for elective care is controlled by waiting lists. However, the model must allow for the perception of a need disappearing with the passage of time as people realise that they do not in fact require treatment.

2.2 The Formal System

In any given area, there will be a range of care facilities - hospital, surgery etc - which will carry out different but overlapping tasks. Although the model should have the capacity to reflect the facilities that already exist, it should also be capable of embodying ways of providing a service that are not currently

available.

The discussion in this section therefore is based round the five core functions that any emergency care system must provide, rather than the characteristics of the currently available forms of provision. These are:

- provision of advice;
- decision-making as to where a potential user of the formal care system should go - this activity is termed *routing*;
- decision-making as to what treatment is required;
- transport provision, and,
- delivery of treatment.

In the NHS as it stands, these activities are closely linked with particular institutions, eg ambulances provide transport, hospitals provide treatment and so on. For our purposes however it seems important to abstract from the specific 'bundles' of activities which each current institution provides, so as to allow for new 'bundles' emerging and for new institutions being developed which focus on only one of these. For example, any provider may in principle offer telephone advice or the provision of telephone advice could be a free-standing service, available from a national or a local source. Similarly, while currently ambulances primarily convey people from home or site of accident to hospital, they could operate in different ways, eg they might take people to a minor injuries unit if their condition was not serious.

<u>Advice</u>

The individual experiencing an urgent event may seek advice as to what to do. In the first instance this can be as simple as an individual considering whether or not his/her perception of a problem requires some further action. Take the case of someone who experiences for the first time a sharp chest pain at 3.00 in the morning and decides it is not serious enough to justify a call to the doctor. In this case the individual is providing the advisory role for him or herself.

Where the individual is not confident enough to do this, they will turn for advice. Advice-seeking may imply travel to the service providing advice; alternatively, advice may be delivered over the telephone, video-link or computer. In our existing system advice is provided by GPs, their receptionists, or other practice staff. Other locations may be the Samaritans or a crisis help line for the mentally ill, or in some cases, an A&E department. Equally a Minor Injuries Unit may provide advice in those areas where it is not competent to provide treatment.

The essence of this activity is that it will influence whether or not the individual further pursues the problem perceived as an urgent event. It will potentially change the action which the individual would have taken in the absence of such advice: for example if telephone advice is available, then the individual may not go to the A&E department 'just to be on the safe side'. In other words, the availability of advice will change - either way - the level of perceived risk the individual associates with any particular set of choices.

Thus, for example, the person with chest pain may believe that this is unlikely to be a serious heart condition. He or she makes a cup of tea, considers the possible options, and decides a good night's sleep will set the situation right. Implicit in this decision is an estimate that the risks of not seeking care are low.

The individual may go one step further and turn to a friend for reassurance, with a similar outcome. The person may contact some form of GP service which also gives the reassurance that there is unlikely to be a major problem. In this case the individual's implicit estimate of his or her own risk profile may be altered, ie the doctor convinces them that the chances of a major problem are even more remote.

Routing

At its most basic the decision of the individual to go back to bed is a routing decision, where the individual is the routing decision-maker. Within the formal care system, the routing decision may be made at the GP surgery where the GP may decide that hospital admission is required, over the telephone where an ambulance controller may decide that the patient requires a particular hospital facility, or in the waiting room of a major A&E hospital where the triage nurse directs the patient to the appropriate professional within it.

How and by whom the routing activity is carried out is a key element in the formal care system. As things currently stand, individuals mostly route themselves to either the GP surgery or the hospital A&E department in the light of their own views as to what services they require and how easily they may access them. Hence the GP surgery and the A&E department are the key points within the formal care system where routing decisions are made.

This pattern in turn reflects the assumption that routing decisions are primarily for doctors. That assumption is being undermined by several developments. As noted already, within the informal system self-routing decisions can be

influenced by the availability of advice. Within the formal system, both nurses and paramedics may in some circumstances make routing decisions in their own right. For example ambulance crews are normally required to take patients to the nearest hospital and their controllers to send an ambulance to each request for help. But experiments are underway which allow them to take independent decisions which may result in a different route to care than would otherwise be taken, or to make such decisions having contacted a hospital doctor by IT link.

Treatment decision

The treatment decision may be made in many parts of the service, from the GP, the nurse, the general hospital doctor, through to the specialist. The treatment decision is normally closely linked to the actual delivery of care - the doctor or nurse determines what should be done and does it, or hands the patient on to someone in the team who can deal with it. But as with routing, it is sometimes helpful to distinguish the decision-stage from the execution phase so as to allow for new ways of working to be developed which comprise different bundles of activity.

The availability of advice *within* the formal care system may be critical to this. For example, a nurse in a minor injuries unit will be confident of her ability to treat in the vast majority of cases: in others, she will be equally sure that the condition presented to her is outside her competence. In the middle are cases which she feels unsure about but which she may be able to resolve if suitably supported by access through IT to advice. With this advice the number of cases she routes to other facilities will be reduced.

Within the hospital A&E department, the decision as to what should be done may involve routing the patient to another place of care which does not require the facilities of the hospital. This way of working remains rare, but some hospitals are beginning to develop triage protocols which provide for the option of 'treatment at home' by a district nurse or some other community-based professional.

Transport

When an urgent event arises there is often a need for transport, either of the individual to the formal care system or for an element of the system to the individual. Thus, in the example above, a GP may visit the individual's home in order to advise the person who perceives a potential emergency problem. Alternatively the individual may be transported to a location where advice will be

delivered. This may be a GP's surgery, a Minor Injuries Unit or a hospital A&E department. Modes of transport may comprise the individual's car, a taxi or an ambulance.

The mode of transport used will determine whether other activities of the emergency care system can be carried out at the same time. For example, an ambulance may be able to deliver life-saving services on the way to the intended location, if the need arose unexpectedly. Alternatively it may bring someone - say a nurse - who can provide care in the home for an indefinite period.

Some care providers are mobile and others not, eg it is rare in the UK for hospital-based services to move to the patient or for ambulances to carry staff more experienced than paramedics. Such combinations are found in other countries and may in some circumstances have a role here.

We make the following distinctions between mobile units, mobile services and transport:

mobile unit: a conveyance equipped to provide some treatment (by implication it can contain both patients and health care professionals);
mobile service: a limited treatment facility that can be moved to a number of sites (including the site of the urgent event) and once there can provide some treatment.
transport: any conveyance (vehicle) that is not equipped to provide treatment. It can convey one or more patients and/or health care professionals.

Treatment delivery

This final attribute of the formal care system, and often the most visible, is the actual delivery of treatment. This may vary from the application of a sticking plaster to a cut knee to open heart surgery, and anything in between. The important point is for the individual to have access to the level of treatment within a timescale which allows an acceptable chance of a successful outcome.

As the earlier discussion implies, care may be delivered in different locations and by different staff. The range of work done in an A&E department, minor injuries unit etc can vary a great deal. The model must be able to allow for this variation and for the creation of new treatment sites or new combinations of treatment and other functions, of a kind not currently available.

2.3 Measures of outcome

So far we have discussed the type of events and individuals which the system may deal with, and the specification of the provider system itself. A final but fundamental question is how the success of the system, or parts of it, is identified. Several outcomes might be considered including:

- activity, eg number of care episodes at each site, overall and within each time period;
- crude measures of outcome such as death;
- travel times/cost and distances;
- delays or waits which it may be difficult to estimate directly;
- resources employed including amount of slack or, put another way, the extent of utilisation.

The key factor determining whether the system as a whole is working well is our attitude to risk. The output of the system as a whole is some level or index of security. Measuring this involves some explicit balance of risk against cost.

Any system must balance the small chance that the individual is about to suffer a fatal coronary at 3.00 in the morning (but one which is treatable in precisely the right circumstances, another quite large assumption) against the impossibility of treating all emergency events in this way, with the resource costs which that would imply. It is not merely a question of financial resources, it may be impossible within current technologies to provide an almost risk-proof service even if it were considered desirable.

Eventually it may be possible to derive a single index of risk for the system. In the initial development of the model, as shown in the next section, we have concentrated on more conventional measures of output, ie activity levels, delays etc.

3 Specifying a Simulation Model

The previous section discussed a whole systems approach to the issue of perceived needs for emergency care and the responses to these. We now describe how that model is mapped into a prototype simulation of how people react to what we have termed 'urgent' events, and the way in which this gives rise to demands on health care services. First, however, we outline how the simulation system works.

3.1 How the simulation system works

The computer system used for the prototype is a software system for simulating models of complex socio-economic systems. It models each person experiencing an urgent event as an individual computer process making its own decisions depending on its current circumstances.

The simulation ensures a formalism to the systems approach. The model has to be complete in order to do a computer simulation. This does not mean that every aspect needs to be known in fine detail but if, for example, there is no model of how patients leave A&E then that facility will soon become clogged up.

The type of simulation used does not set out to model the lengths of queues – if the system being modelled is such that individuals have to wait before they can satisfy some need then a queue will emerge as simulated time progresses.

The computer system uses sophisticated language technology for mapping of models of the real-world onto a calculus of communicating processes (Milner, 1989; Bryan-Jones and Tansley, 1997). This means that a simulation is many autonomous computational processes each proceeding independently of, but able to communicate with, the others. The language is used to write *scripts* that can be easily validated by those with knowledge about the system being modelled. This is termed *clinical programming* and *management programming*. There are scripts for clinical protocols, management policies, effects of co-morbidity, demographics and other personal circumstances, descriptions of what people (patients and staff) do, their perceptions, decisions and actions. Using the system is an iterative process: simulate, compare output with reality, revise the model and its programmed version, simulate.

3.2 Simulating the environment

In broad terms, what happens in the simulation can be thought of as the creation of a succession of virtual people who have (or think they have) just experienced a virtual urgent event. As time progresses, what these virtual people do depends on their attributes and their circumstances. They remain in the simulation while they are 'managing', 'seeking advice' or 'seeking treatment'.

Simulating a virtual population experiencing urgent events

The simulation generates virtual 'urgent events' happening to virtual people. The generation process has access to a lot of information about how, when and what sort of urgent events occur, their signs and symptoms and likely progress. This includes how incidence and progress relate to age and other demographics, how they vary geographically, in different seasons and at different times of day, how they are affected by the circumstances of the event and the subsequent behaviour of the individual, other people and the health services. The generation process also has access to information about types of people, their knowledge of health care and health care services, and factors that are likely to affect their decision-making.

Each simulated individual has *attributes* containing personal characteristics and medical history (as far as they are relevant to the reaction to the event and subsequent progress), signs and symptoms and information about the incident and factors that are likely to affect their subsequent progress. Examples are self-diagnosis ability (medical knowledge in general or knowledge of their own condition), personality, socio-economic status. The factors may be correlated - eg self-diagnosis ability may be higher if there is a long medical history and the event is not unexpected. Attributes may be modified: for example a particular description may be 'gradually get better unless anxious' and if the virtual patient is not anxious, they will show a gradual improvement in the simulation. If, on the other hand, sources of reassurance such as an advice line are not available, then the patient may become anxious and their condition deteriorates.

The knowledge base is compiled from several sources, ranging from population statistics and evidence-based knowledge to consensus and educated reasonable assumptions. The prototype has a knowledge base of over 20 different categories of urgent event, and this can be extended. A user of the simulation has access to this knowledge base through the generating process and so has a lot of flexibility in the specification of the mix of individuals and urgent events to be simulated.

The prototype works on an hourly beat, ie every simulated hour it creates a new set of people who have urgent events, and adds them to those still in the system. Typically, 10 per cent of the urgent events in the simulation have just occurred and 90 per cent have been progressing for at least an hour.

Simulating clinical conditions and severity – 'event profiles'

An individual's progress after experiencing an urgent event can follow various patterns. It is modelled as multi-dimensional variables representing several signs, symptoms, measures and outcome indicators. The shapes are related to personal factors, the clinical conditions and clinical intervention; some examples are shown in Figure 1. We term these *event profiles*.

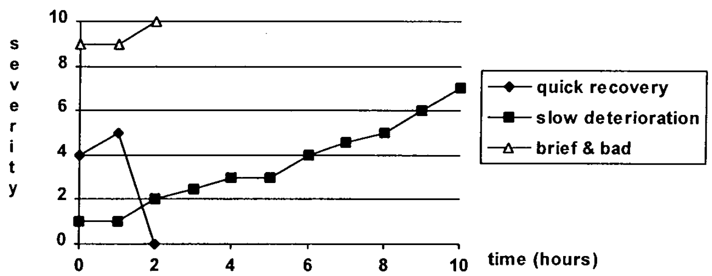

Figure 1: Shape of some event profiles

Event profiles may be based on strong evidence-based knowledge or reasonable, but uncorroborated, assumptions. The information about event profiles is part of the simulation knowledge-base; it can be seen as a 'library' of descriptions of the progression of urgent events and patient groups. These descriptions are stored in the form of data tables and scripts, and can be read or browsed by users.

The prototype has a simple event profile for each simulated urgent event represented by a 'severity score' on a scale from 0 to 10. As time progresses this severity score may change - showing 'improvement' or 'deterioration' - and it may also be affected by intervention. A probability of death can be associated with high scores; the prototype model assumes that this is zero unless the individual reaches the most severe score (10). To model different life-threatening situations that could be averted, we associate a different probability of dying if the severity remains at 10 for one, two or three hours without medical treatment.

Simulating actual and perceived needs

The model described in the previous section concentrates on the perceived risk and the responses to it. The model recognises that risk assessment is the fundamental activity which is going on. It also recognises that, with the degree of uncertainty which exists in the emergency care system, there will be occasions when it fails the patient. The skill is to design the service to minimise the risks of this happening.

The simulation needs to be able to make this explicit and therefore we include a model of the real risk, as well as modelling perceptions. These are modelled through the event profiles.

There are several possible views of an individual's clinical state. The patient has a view, there are one or more clinician views (to cover the different professions or individual interpretations) and the actual state. Table 1 indicates how these may affect the individual and their progress in the health care system. In the prototype there is a single clinician view and there is relatively little distinction made between the personal and clinical views and the actual severities.

Table 1: Simulating different views of an individual's event profile

Whose view	*number*	*Description*	*Influence*
patient	1 view	The individual's perception of the severity of the event.	Affects how and when the individual acts.
clinician	several views	These represent different professions or individual clinicians' views of severity.	Affects how the components of the formal system respond
actual	1 view	This is the actual clinical state.	Affects the probable outcome of various actions.

How individuals react to their perception of need

Each individual is in one of three states: they are seeking treatment, seeking advice or 'managing'. This continues until they leave the simulation for one of three reasons: they have stopped using the formal health care system, feel recovered or have died.

Individuals seek treatment based on some personal decision mechanism, possibly after taking advice. In the prototype, individuals seek treatment if their personal decision score exceeds a defined threshold. The score combines severity (own perception), age, a knowledge factor as shown by their self-diagnosis ability, and

a personality factor reflecting their self-reliance/likelihood of seeking help. The threshold is a fixed quantity, which can be modified by the user, but is the same for all individuals and all types of urgent event. However, as with other parts of the simulation, it can become a more detailed model of the reality where the threshold is much more dynamic, and will change according to what is happening elsewhere. The default setting of the threshold results in about 1 in 25 (4%) of those who experience an urgent event deciding to approach an agency and hence potentially entering a component of the emergency care system. Figure 2 shows the potential demand on the formal system in a simulated experiment.

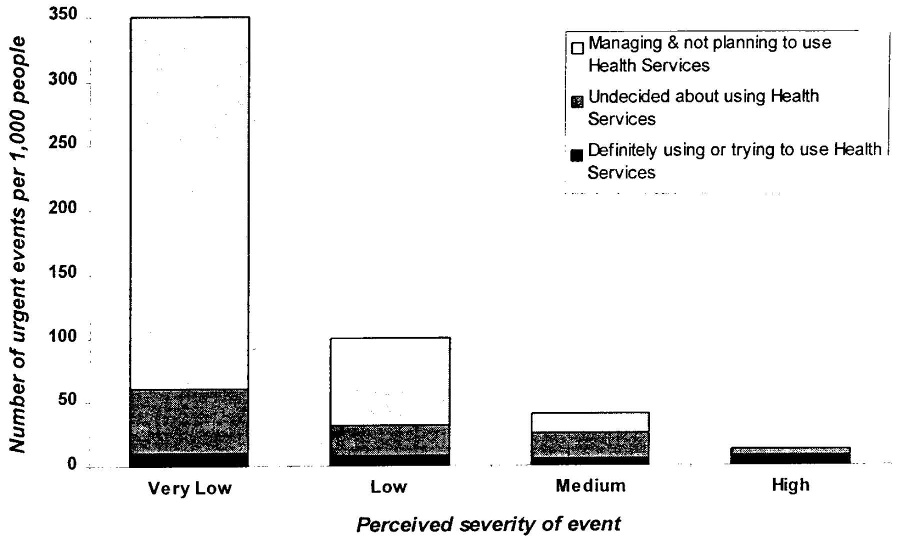

Figure 2: Example of the simulated potential demand on formal system

In the prototype, if an individual seeks treatment but cannot get it, then they just wait (manage) until either the relevant treatment facility is available, or changes in their perceived event profile cause them to seek treatment elsewhere or to stop seeking treatment.

In the prototype, individuals are more likely to seek advice if they are in socio-economic group A, B or C1 and their knowledge factor (self-diagnosis ability) is less than a critical value, (this is initially set at 2, but can be changed by the user). The decision to seek advice rather than treatment may be affected by the expected travel times and perceived waiting times at the different providers. If an individual seeks advice then they may or may not act on the advice that they receive. In the prototype, if an individual seeks advice then they act (or try to act) on the advice immediately (ie within one hour). Thus, if advised to seek

treatment, they try to go to a health care/treatment facility.

3.3 Simulating the response from the formal health care system

The simulation includes all health care services which can deal with, or advise about, urgent events, plus services that it is felt the public might approach in connection with urgent events. We describe next how it deals with the five functions set out above which any emergency care system should address.

Consultative or advisory

The prototype has a range of 'health care facilities', providing treatment and/or advice. These include GP, Out-of-hours Services, Minor Injuries Unit, A&E, Specialist Centre, Mobile Treatment Facility and an Advice Centre. Each has scheduled opening hours and a capacity which may vary according to circumstances. An individual can access directly all facilities except specialists.

All care facilities can give advice – if clinicians are 'present' then they advise (using the information available to them); otherwise the advice is based on the individual's perception of illness severity.

Routing decision-maker

In the prototype, if an individual decides to seek advice then they try the Advice Centre, which responds immediately as it is assumed to have unlimited resources.

When an individual decides to seek treatment, they choose a Treatment Facility according to their perceived illness severity score. In the prototype, if the score exceeds 5 they go to A&E, otherwise they call/go to their GP. The prototype routes them to the nearest facility of the type they have chosen to attend, where 'nearest' is a simple measure of time. Each facility can route an individual to a more appropriate facility.

Certain aspects of decision-making while experiencing an urgent event depend on where it occurs and how far this is from the treatment centres/services. In the prototype, this is simulated using a simple map marked into zones (eg Figure 3). The map can show the location of treatment centres, indicate different travel times and can be used to help set the scenarios for simulations.

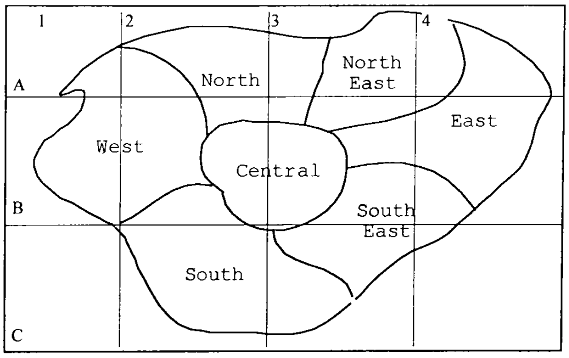

Figure 3: Area map showing zones

Treatment decision-maker

Each of the care facilities modelled can treat a specified range of illness severity, as determined by the simulated clinicians.

Transport provider

The simulation has a knowledge base describing different types of mobile facility and transport. Mobile services and mobile units can provide a range of treatment facilities and can be summoned by appropriate services in the system. Also, the availability of transport or waiting times may vary according to location, time of the day or what is happening elsewhere in the model.

Travel times will depend on start-location, destination and time of day. The actual specification will depend on the detail of the model. Separate timings may be required for different types of vehicle.

The prototype has a very simple model of transport. How the person moves or is moved between the site of the event, ie where they are when they contact the care facilities, or how they get to a particular care facility, is represented by the time taken for the type of transport used to move between zones. In the prototype, each mobile facility is modelled as available or not within each location.

Treatment deliverer

The prototype has a very simple, idealised model of treatment delivery. Once the patient has arrived at an appropriate treatment provider it is assumed that this treatment is given correctly and efficiently, the patient recovers and leaves the simulation. It is straightforward to model a proportion of other outcomes based

on severity of condition.

There may, however, be a delay in delivering the treatment and this is modelled through the capacity of the deliverer. In the prototype, there is unlimited capacity when the facility is open, but the user can restrict this. (Remember that delays and waiting on entering the formal system are modelled in the consultative and advisory functions.)

3.4 Using the simulation

Prototyping contributed in two ways to the development of the model described in Section 2. First, computer specification requires the model to be complete, even if it is simple - forcing a rigour to the modelling that could otherwise be avoided. Second, the simulation results provide useful feedback about the accuracy and appropriateness of the model.

The prototype can run several scenarios simultaneously, displaying performance indicators (outcomes, waiting times, costs, etc) on daily charts (totals, ranges, Z-charts), distributions in time or by value. Patient histories and staff/resource usage can be saved in csv-files for transfer to spreadsheet, database or statistical packages for analysis.

The system is particularly suited for designed simulation experiments. It uses a concept of *foreground* and *background* scenarios. Foreground scenarios relate to factors/options that are being investigated – they are the factors that can be changed. Background scenarios are examples of situations that may occur – perhaps they have occurred in the past and caused problems (or were luckily averted but could easily happen again). The aim of the experimentation is to make the foreground comparisons against all the background scenarios. The experiments can be set up using formal statistics (complete or fractional factorials, maybe drawing on Taguchi simplifications) or can be ad-hoc – but, of course, the former is recommended.

The system can be used for service design. Simulations can confirm or question the expected results of different configurations. Cost/performance implications can be examined and discussed in detail (before disruption to existing services). Sensible, drastic, crazy or conservative ideas can all be tried out – conflicts can be resolved, better specifications encouraged.

The system can be used to demonstrate a benchmark and to compare current and intended formal services against that benchmark. The system can be used to

develop a related set of process control indicators - warning and action limits and measures of variation.

3.5 Extending the simulation model

The prototype can be quickly and easily extended to become a more detailed simulation of the complexities of the framework described in Section 2. We give some examples here.

- Information can be added to the knowledge-bases and the generating process and protocols (clinical and management) can make use of it. For example, the generating process can be extended (by programming) to cope with more types of urgent event, multiple events and epidemics. Populations can be specified as a health authority, a rural/urban or rural population etc. An individual's response can take account of geography, location, season and time of day.
- The simulation of the effect of different perceptions of need can be more subtle. For example, the accuracy of the patient's view depends on the individual's level of knowledge and experience and their personality (self-reliant vs. worrier). The accuracy of the clinical views depends on the type of urgent event, the level of clinical experience and expertise of those involved and the level of (evidence-based) clinical knowledge of such urgent events. Professional and experiential differences can be included.
- The simulation of the individual's decision-making can be more detailed. The probability of seeking advice or treatment can be related to the event, its duration and the perceived current severity, expected travel times, perceived waiting times and so on. It can also involve feedback, eg time since event, time of day, effects of people's learning and access to improved information, no response from the Advice Agency, etc. The simulation of the effect of location and time of day on choice of transport and journey time can be more sophisticated.

4 Conclusion

This approach to modelling in health services is relatively new. It integrates a whole systems approach with simulation software. It provides a formalism to systems thinking and opportunity to evaluate different design options to see how adaptable they are to future scenarios. We term this *adaptive whole-system design.*

The prototype has been an integral part of the development and documentation of a whole systems model of urgent events and associated formal health services. The computer requirement for precise specification sets a new discipline of rigour in the whole systems model. The model must be completely specified – even if some parts are very simply described. Thus, this approach quickly indicates where knowledge is weak or lacking and hence shows clear requirements for further detailed research. It has already contributed to our understanding, and has the potential to enable and encourage a wider mutual learning about this complex system.

The prototype shows how particular policy options can be tested within a well-defined assumptive framework before actual implementation takes place.

The model is capable of almost unlimited extension to take into account more complexity of the real world. The next stage is to test its potential in a real situation. That stage will, we hope, begin soon.

References

S Kendrick, The pattern of increase in emergency hospital admissions in Scotland, *Health Bulletin*, March 1996, S4(2), pp 169-183
N Edwards and V Warneke, In the Fast Lane, *Health Service Journal*, 1994, 8 December pp130-132.
R Milner, *Communication and Concurrency*, Prentice Hall Europe, Hemel Hempstead, 1989
J Bryan-Jones and J Tansley, *EpiScript User's Manual*, An Teallach Limited, 1997