

Examining the feasibility of using a modelling tool to assess resilience across a health care system and assist with decisions concerning service reconfiguration

Abstract

Changes in medical practice, demographic shifts and financial pressures are all examples of factors that may contribute to demand for periodic changes in the configuration of health services. When reconfiguring a service, health planners often take into account projected demand for services, patient access criteria and budgetary constraints (amongst other things), but typically give little consideration regarding its resilience to deliver services during and after external disruptions to its capability to deliver. In this paper we discuss a study conducted in response to a direct request from the National Health Service (NHS) Resilience Project within the Department of Health to explore the feasibility of assessing resilience across local services within the NHS and developing a computer software tool to assess resilience of different service reconfigurations. We give an account of the modelling process used, including the analytical framework we developed using both optimisation and heuristic methods, and an illustrative example of usage of a prototype software tool. We also highlight the key lessons that emerged during this project, which may be helpful to OR analysts working on similar projects regarding resilience in the public sector.

Keywords

Practice of OR, Health service, Optimization.

Introduction

Emergency preparedness can help reduce the impact to society and the economy of major disruptions such as fuel shortages, an influenza pandemic or widespread flooding. Preparing for such events may comprise a range of measures and is often required to be co-ordinated across local, regional, national and sometimes international borders. In this paper, we focus on emergency preparedness in health care and, in particular, those aspects that are related to resilience. In the context of this work and as defined by the project's client, we use the term resilience to mean the

capability of a health system to mitigate the impact of major external disruptions on its ability to meet the needs of the population during the disruption. Considerations of health system resilience may include strategic decisions, such as the allocation of health service provision across different sites, as well as operational decisions, such as the design of robust stock management for essential health care supplies. In this work we focus on the former.

Periodic alterations in the configuration of health services arise as a result of political cycles, changes in medical practice, demographic shifts and financial pressures amongst other things. The decision-making behind reconfiguration is complicated and multifaceted, with health planners taking into account factors such as budgetary constraints, projected demand for services, the accessibility of services to patients, economies of scale and quality of service provision (Imison, 2010; Fulop *et al.*, 2010). The configuration of a health system can affect its resilience. For example, reconfiguration often involves the concentration of services to enhance safety, effectiveness and efficiency, but this might result in a system with key services available at fewer sites which may be more liable to disruption. Yet it is not common for resilience to be taken into account explicitly in decisions concerning service reconfiguration within the English National Health Service (NHS). The desire to routinely and systematically include considerations of resilience within planning motivated the NHS Resilience Project within the UK Department of Health to instigate the research presented here.

In this paper, we give an account of an Operational Research (OR) project conducted in response to a specific request from the Department of Health to explore the feasibility of assessing resilience across local health services and developing a computer software tool to assist with decisions concerning service reconfiguration in the NHS in England. The aim of the work was to investigate the potential role of OR in assessing resilience at a local-regional level in England and to explore the feasibility of developing a computer software tool to assist with decisions concerning service reconfiguration. We note that this request was made within the context of a broader UK-wide program on resilience (Cabinet Office, 2013) and other initiatives relating to the NHS being conducted by the Department of Health (e.g. Department of Health, 2008).

In the next section we describe in more detail the background of this work and give a brief summary of some relevant literature in the field. We then give a detailed account of the project, including an outline of an analytical framework devised, a model and prototype software tool developed to facilitate engagement with the envisaged end users, and a brief example of illustrating the use of the tool. Finally, we summarise the project outcomes and discuss the insights generated by the project including the challenges faced and the key lessons learned.

Background

In recent years there has been a number of relevant editorial, review and position papers on the role of OR in emergency preparedness and emergency response management (Altay and Green, 2006; Brandeau et al., 2009; Green and Kolesar, 2004; Larson, 2004; Larson, 2006; Simpson and Hancock, 2009; Wein et al., 2009; Wright et al., 2006). The problems tackled cover a wide variety of application areas. Up until the 1990s, the majority of work in the field focused mainly on ‘routine emergencies’ within the context of established emergency services such as fire, police and ambulance (Simpson and Hancock, 2009). Research was largely concerned with determining the optimal number, location and allocation of response units within municipal services, as exemplified by the large body of influential work emanating from the New York City - RAND Institute (Green and Kolesar, 2004; Larson, 2002).

More recently, there has been a shift in focus away from routine emergencies towards larger scale ‘disaster emergencies’, which occur “when resources become stressed, when non-standard procedures must be implemented to save life or when special authorities must be invoked to manage the event” (Altay and Green, 2006). Much of this work focuses on applications in security (Larson, 2004; Wein et al., 2009; Wright et al., 2006), the health care component of which is mainly concerned with the response to bioterrorist threats such as deliberate introduction of anthrax or smallpox (Hupert et al., 2002; Lee et al., 2009; Wein et al., 2003; Zaric et al., 2008). Other published examples of OR applied specifically to health care within a disaster emergency context, of which there are relatively few, include planning for emergency mass dispensing and vaccination clinics – in a pandemic influenza or otherwise (Aaby

et al., 2006), the routing of ambulances to hospitals in a disaster (Jotshi *et al.*, 2009) and the design of interventions in a pandemic influenza (Eichner *et al.*, 2007).

An extensive selection of OR techniques has been used by researchers in the area of emergency preparedness. The most common appear to be mathematical programming including heuristics (e.g. Rolland *et al.*, 2010), probability and statistics, and simulation (Green and Kolesar, 2004; Simpson and Hancock, 2009). More specifically in health care, discrete event simulation has been applied to the problem of designing antibiotic and vaccination distribution centres in the case of bioterrorist attacks or pandemic influenza (Aaby *et al.*, 2006; Hupert *et al.*, 2002; Lee *et al.*, 2009), whilst compartmental modelling and queuing theory have been used to compare various emergency responses in the event of an anthrax attack (Wein *et al.*, 2003; Zaric *et al.*, 2008). Tufekci and Wallace (1998) point out a lack of information and decision support systems in the field, whilst Simpson and Hancock (2009) note that OR studies within disaster emergency preparedness have had relatively little influence on policy or practice, highlighting the need for positive engagement with responders and policy makers in future research.

To the best of our knowledge, there are no OR studies reported in the literature that address the problem of health system resilience. Whilst we have not conducted a systematic review of the area, a search of the PUBMED, Web of Science and INFORMS databases using the keywords “resilience” and “health system” or “health service” did not return any relevant articles (although we acknowledge that these search criterion may have excluded other related keywords used by authors). A similar outcome was obtained by back-referencing the recently published review articles in OR and emergency preparedness or disaster planning (Brandeau *et al.*, 2009; Green and Kolesar, 2004; Simpson and Hancock, 2009). The relatively small body of OR work applied to health care in disaster emergencies is generally focused on a specific type of disaster and/or a narrow range of emergency responses affecting a small part of the health care system. The work reported in this paper is more generic, motivated by our aim of scoping a tool for use in aiding decisions concerning service reconfiguration at a strategic level.

Modelling process

MODELLING APPROACH

Given the remit of the project (to explore the feasibility of developing a computer software tool to assess resilience of different service reconfigurations), and the wide scope and complexity of the problem, we decided to formulate a deterministic model in the first instance. This approach was based on the notion that if one cannot build and populate a simple model of service delivery, as required for the purposes of this problem, then doing so for a more complex, stochastic model would be unrealistic.

The intended use of the proposed model is to give an assessment of the impact of a given pattern of disruption to health care resources and infrastructure, including supply chains, on the capability of the health system to respond to the attendant pattern of demand (in terms of number of cases) for certain care services. In doing this, it seems sensible to consider feasible responses to disruption in terms of reallocation of resources. In this work we have chosen to incorporate a range of such responses within the model (using optimisation and heuristic methods to capture response) to take into account inherent uncertainty surrounding the nature of this response.

The schematic diagram in Figure 1 gives an overview of how the proposed modelling tool fits into a broader analytical framework. The specific cause of the disruption under consideration (e.g. pandemic influenza, inland flooding or severe weather) does not explicitly inform the assessment of impact other than by informing, in some process outwith the scope of the model, both the pattern of disruption to health care infrastructure and the pattern of demand for acute care services. We envisage that users would generate disruption scenarios in line with regional and national initiatives for disruption planning such as the Community Risk Registers (see, for example, Warwickshire County Council, 2011) and the National Risk Register for Civil Emergencies in the UK (Cabinet Office, 2008). The pattern of disruption associated with a scenario is taken as an input to the model and is defined by reductions in resources (in the broadest sense) at specific hospitals. In a similar fashion, the pattern of demand associated with a scenario is taken as an input to the model and is defined by demand for health care services at each hospital.




Figure 1: A schematic diagram showing the scope of the model (inside the dashed line) within a broader analytical framework.

MODELLING TOOL

Conceptually, the model represents a stylised health system comprising three basic entities: care services (each a combination of a defined set of health care interventions for a specified group of patients), hospitals and resources (such as staff, clinical environments, utilities, and clinical and non-clinical supplies). Figure 2 illustrates the relationships between these three entities. The ‘requirement’ relationship gives the numerical amounts of each resource required to deliver care services, whilst the ‘availability’ relationship gives the amount of each resource available at each hospital (intrinsically or via a supply chain). The demand for care services (‘demand’) is quantified in terms of the number of cases.




Figure 2: The three conceptual entities (hospital, service and resource) and their relationships as used in the model. Resources are assumed to be intrinsic to a hospital or available via a supply chain (organisations and/or companies that are not under the direct control of the NHS).

Summary of modelling assumptions

The health care system is modelled in simple terms akin to an industrial process, whereby a range of different activities (health care services) take place, each requiring resources of different types. We assume that hospitals may provide a number of care services and that each individual case is managed entirely at a single hospital. Furthermore, we assume that there is no variability in the resources required to deliver a single case of a given service, either between individuals or across hospitals. Disruption to the availability of resources is assumed to be fixed over a single period of interest. We also assume that demand for services is deterministic and that, immediately prior to the period of interest, demand for health services is met. Finally, we assume that resources within each hospital can be re-allocated and that, subject to resource constraints, the level of each type of care service provided can be adjusted.

Resilience planning based on unmet demand

The focus of our analysis is the unmet demand that may occur as a result of disruption to resources or a changed pattern of demand, alone or in combination. In order to estimate unmet demand however, a view has to be taken as to what might constitute a plausible response to disruption on the part of the local health system. Given that we do not know what the precise nature of such a response will be, we formulated a range of possible responses of varying degree of sophistication (using optimisation and heuristic methods) based on discussions with stakeholders and for the specific purposes of this feasibility study. We note that some unmet demand (e.g. elective cases that can safely be delayed) could potentially be deemed less important than others (e.g. emergency cases that cannot safely be delayed). For this reason, the model was designed to allow for an optional weighting to be associated with the perceived importance of each service by the user. We did not address in our research the actual method for allocating these weights.

Optimisation-based response to disruption

The optimisation-based approach, which represents the most sophisticated response to disruption considered, is based on an assumption that resources would be optimally allocated to services within and between hospitals in order to minimise unmet weighted demand. Technically, an integer programming problem is formulated to express the problem of minimising unmet weighted demand by reallocating resources and determining levels of different forms of care service within each hospital. The optimisation is subject to constraints related to the observation that care service activity over the period of interest (in terms of number of cases treated) cannot exceed demand (and must be non-negative) and that such activity is also constrained by the availability of each resource. A mathematical description of the optimisation model is in the Appendix.

Heuristics-based response to disruption

In addition to the integer programming method, three different heuristic algorithms were formulated, each reflecting plausible rules for resource allocation in the face of disruption. The first algorithm represents a ‘business-as-usual’ approach until the first resource constraint is reached. It calculates the service activity under disruption by scaling down the pre-disruption activity by the minimum feasible proportion given the reduced resource availability. The second algorithm considers each resource in turn and scales down the caseload delivery by an appropriate factor to reflect the reduction in resource availability. Finally, the third algorithm attempts to prioritise the highest priority care services (ranked according to user-defined weights) as far as possible by, potentially, redirecting resources away from services deemed to be of lower priority. A technical description of all three algorithms is in the Appendix.

Suspension of non-emergency care services

In the description above we have assumed that service managers will attempt as far as possible to maintain the delivery of all types of care services during a disruption (with some services potentially having higher priority). An alternative response, which was included in the modelling tool, may be to temporarily suspend low priority care services in times of disruption to free up resources to deal with more urgent or higher priority care services.

Prototype software tool

A working version of a prototype software tool was developed to facilitate the engagement with potential end users and project stakeholders and the communication of the proposed model's structure and data requirements. It was also used to test the computational feasibility of the envisaged mathematical framework by enabling the construction of illustrative examples (see below for example). In the prototype, the user can enter realistic test data in terms of number of services, resources, hospitals and the supply chain, as well as an optional weighting associated with the perceived importance of each service. The relationship between services and resources and the availability of resources at each hospital are also user-defined. Scenarios are specified as changes to the quantity of resources available at each hospital and the post-disruption demand for services associated with each hospital. The prototype was developed using MS Excel, making use of Solver to implement the optimisation problem and Visual Basic for Applications (VBA) routines to code the heuristic algorithms.

ILLUSTRATIVE TOOL USAGE

As part of the process of engaging with project stakeholders, we populated the prototype software using an illustrative example involving a hypothetical local health system under different service configurations. We used the example to showcase the analysis of the effects on the health system of resource reductions of a type and magnitude thought to be plausible by the sponsors of this study and of interest to them within the remit of this work. The model parameters were estimated, in part, by publicly available data regarding the English NHS (where possible) and with other estimates based on expert opinion. We stress that the example was developed for illustrative purposes only and is not intended to represent any real health care system. A fully calibrated case study would potentially consider a larger range of services and resources (for example, a richer and more representative range of staff inputs and some attempt to capture the range of equipment and drugs used in care provision) as well as more detailed data describing the demand.

The example was designed to reflect some stylised aspects of acute care service provision in an area with a population of about 500,000 people. We considered three

service configuration scenarios informed by a well-documented tendency towards centralisation in healthcare (Academy of Medical Royal Colleges, 2007; King's Fund, 2007; Pollock et al., 1999). The first was the 'status quo' scenario under which all services take place at all hospitals (Figure 3a). In the remaining two scenarios paediatric and paediatric and maternity care services were centralised at a small number of the hospitals in the local health system (Figure 3b and 3c).

The reductions to the resources at each hospital were of a type and magnitude that might reasonably be associated with widespread inland flooding within the catchment area of the health system as outlined in national and regional emergency preparedness planning scenarios. We considered three levels of resource disruption associated with the flooding, corresponding to mild, moderate or severe disruption (Table 1).

Table 1: Disruption levels and corresponding resource reduction

Disruption scenario	Disruption to availability of beds	Disruption to availability of staff	Disruption to own-source resources	Disruption to supply-chain resources
mild	0%	30%	10%	20%
moderate	0%	50%	20%	40%
severe	20%	70%	30%	70%

The adoption of these resource disruptions was based on the following assumptions:

- Flooding is more likely to disrupt the supply chain provision of resources than on-site resource provision due to disruptions in the transport infrastructure.
- Flooding is unlikely to affect the availability of beds unless it is severe and directly affects hospital infrastructure.
- Flooding is likely to cause a relatively large scale disruption to staff availability via its effects on the transport infrastructure.




Figure 3: Schematic representation of the example scenarios of service reconfiguration and disruption considered.

We assumed that flooding causes resource disruption at three of the medium-sized hospitals, two of which coincide with hospitals at which maternity and paediatric services are centralised (the third hospital at which maternity services are centralised was assumed not to have been affected), Figure 3d.

The resilience of acute care service provision under differing levels of service centralisation was analysed in terms of unmet demand for a range of disruptions (of differing severity) that affect resource availability but maintain fixed levels of demand. The model inputs and the particulars of the disruption and centralisation scenarios are outlined briefly in Table 1 and Table 2; further details of the numerical inputs used, along with the assumptions used to arrive at these inputs, are contained in the Supplement.

Table 2: Model inputs and description of the disruption and centralisation scenarios of the illustrative example (further details of the numerical inputs used, along with the assumptions used to arrive at these inputs, are contained in the Supplement)

Model inputs	Description
<i>Care services</i>	Eight care services are offered within hospitals of the local health system, namely: Paediatric, Maternity, Surgery and Acute Medicine, each subdivided into emergency and non-emergency care services.
<i>Hospitals</i>	Ten hospitals of varying size are included in the example. It is assumed that of these hospitals, two are large, five are medium-sized and three are small, in terms of demand for various services (see Supplement). The total demand for services was considered to remain unchanged under disruption.
<i>Resource availability</i>	The resources available to each hospital prior to disruption were calculated based on the demand for services, the resource requirement to deliver each service and the assumption of zero unmet demand prior to disruption. The Supplement contains further details on the calculations involved.
<i>Configuration options</i>	In addition to the ‘status-quo’ option already discussed, we considered two possible reconfiguration options (see Figure 3): Option 1 - Centralisation of paediatric services. Paediatric services are concentrated at two medium-sized hospitals, with total demand for these services unchanged from the status quo but distributed equally between the two specialist paediatric hospitals. Option 2 - Centralisation of paediatric and maternity services. Paediatric services are concentrated as in option 1; maternity services are concentrated at three medium-sized hospitals, with total demand for these services unchanged from the status quo but distributed equally between the three specialist maternity hospitals.
<i>Scenarios</i>	The reductions to the resources at each hospital were of a type and magnitude that might reasonably be associated with widespread inland flooding within the catchment area of the health system as outlined in national and regional emergency preparedness planning scenarios and agreed by project stakeholders (Figure 3). We considered three levels of resource disruption associated with the flooding, corresponding to mild, moderate or severe disruption. The details of the numerical effects on resource availability are summarised in Table 1 and the Supplement.

Results of illustrative example

The model was run for each of the three service configurations and each of the disruption levels (mild, moderate and severe). The outputs are summarised in Tables 3-4, where unmet demand is reported for the optimisation and three heuristic

modelling approaches described previously. Table 3 summarises results where it is assumed that both emergency and non-emergency services are retained whilst Table 4 summarises results under the suspension of non-emergency services (see relevant subsection above).

The results indicate that in resilience terms, centralisation may reduce the ability of the health service to meet demand for health care. Unsurprisingly, this is especially true if disruptions occur at hospitals where specialties serving a wide catchment area are located. As expected, more severe disruptions are associated with higher unmet demand across the range of system responses. The scale of unmet demand varies across the range of models, with its smallest value for the optimisation and largest outputs corresponding to the single-step scaling algorithm.

Table 3: The impact of flooding on services, expressed as percentage of unmet demand for all services (retaining both emergency and non-emergency services).

Disruption severity	Configuration	Optimi-	Heuristics response		
		sation response	Resource-based	Service-based	Single step
Mild	No centralisation	5%	8%	8%	8%
	Paediatrics centralised	7%	10%	9%	10%
	Maternity and paediatrics centralised	8%	12%	10%	12%
Moderate	No centralisation	10%	13%	13%	14%
	Paediatrics centralised	12%	17%	15%	17%
	Maternity and paediatrics centralised	14%	20%	18%	20%
Severe	No centralisation	15%	19%	18%	19%
	Paediatrics centralised	19%	24%	23%	24%
	Maternity and paediatrics centralised	23%	28%	26%	28%

Table 4: The impact of flooding on services, expressed as percentage of unmet demand for all emergency services, with non-emergency services suspended.

Disruption severity	Configuration	Optimisation response	Heuristics response		
			Resource-based	Service-based	Single step
Mild	No centralisation	0.4%	4%	0.4%	8%
	Paediatrics centralised	0.4%	5%	0.4%	10%
	Maternity and paediatrics centralised	0.9%	7%	0.9%	12%
Moderate	No centralisation	5%	8%	5%	14%
	Paediatrics centralised	5%	9%	5%	16%
	Maternity and paediatrics centralised	8%	13%	8%	20%
Severe	No centralisation	14%	14%	14%	19%
	Paediatrics centralised	15%	18%	15%	22%
	Maternity and paediatrics centralised	20%	24%	20%	28%

Engagement with stakeholders

The stakeholders in this project can be separated into two groups, namely the “client” who defined the project remit and commissioned the work, and the “envisaged end users” of any potential software tool. The former was the UK Department of Health and the latter are, broadly speaking, national and regional emergency planners. At the time of conducting the project, the client defined the primary envisaged end users to be the emergency planning leads of the 10 UK Strategic Health Authorities (SHA EP leads).

Throughout the project we held regular discussions with the client regarding the remit of the work, development of a suitable modelling approach, engagement with end users, illustrative prototype usage and the lessons learned from the feasibility study. We used the model and accompanying prototype software tool and illustrative example to raise awareness of the complex issues involved and to provide a focus for these discussions. We also used them to engage with the SHA EP leads in considering data requirements and their availability within the health service, and in seeking opinion on issues around suitability and usability of tool.

The early stages of the modelling approach were informed by interviews with two SHA EP leads. Subsequently, a meeting was held with all of the SHA EP leads in which we addressed the following questions: Would the envisaged output of the tool be useful to you? What level of detail would such output need to have to be useful? Would it help you contribute to the decision-making process? What level of confidence in the tool would you need to have? Supplementary questionnaires regarding the availability of parameters that would need to be estimated based on local knowledge were completed by SHA EP leads as well as a consultant in General and Geriatric Medicine, a consultant in Cardio-Thoracic Surgery and a senior operations manager of a large teaching hospital. The outcomes of these consultations and the lessons learnt informed the final decision taken by the client.

Discussion

This paper gives an account of an Operational Research (OR) project conducted in response to a specific request from the National Health Service (NHS) Resilience Project within the Department of Health to explore the feasibility of assessing resilience across local health services and developing a computer software tool to assist with decisions concerning service reconfiguration in the NHS in England.

We developed a generic analytical framework for modelling the impact on care service activity of potential disruptions to a health system and, as a means of communicating with stakeholders and to test the feasibility of the suggested model and subsequent software tool, we constructed a prototype tool and an illustrative example that involves a hypothetical local health system under different service configurations. We deliberately adopted a modelling approach that led to simple and static model of service delivery, as opposed to a more complex, stochastic model that would have allowed the capturing of notions of service variability, on the assumption that if it is infeasible to build and populate the former, then it would also be infeasible to build the latter. System dynamics, in particular, was considered and rejected on the basis of the need for a simple, static approach and the perceived lack of added value to the modelling in this particular study by including feedback loops.

We note that our approach does not address the capability of the system to return to routine activity, which may involve dealing with a backlog of routine activity resulting from resources being diverted to deal with the disruption (e.g. increased waiting lists for elective care). Nor does our approach allow for disruptions to health systems with existing waiting lists (we use a simplifying assumption that, immediately prior to a disruption, demand for services is met), and we note that the long terms effects on such services could be substantial even if the immediate emergency response is adequate.

Both the mathematical model and accompanying prototype software tool differ substantially from other published modelling work in the field of emergency preparedness, not least because we addressed issues of strategic decisions in health system planning as opposed to operational issues arising from disruptions (Green and Kolesar, 2004; Simpson and Hancock, 2009). Whilst the benefits of computer simulation methods have been exploited in a number of context-specific applications within the field (Aaby et al., 2006; Hupert et al., 2002), our choice of modelling methodology arising from the envisaged use of the tool, has the benefit of lending itself more readily to iterative use as a planning tool and ultimately free distribution to end users.

A number of learning points have emerged during this project, both from the research and development conducted in relation to the modelling tool and from our engagement with potential end users and other health professionals. From a health system's perspective, there seems to be very limited detailed knowledge on the number and quantity of services and supplies that are outsourced or provided through external supply chains (organisations and/or companies that are not under the direct control of the NHS). In addition to having an impact on the calibration of the tool, it raises the related question of the actual resilience of the NHS supply chain.

The envisaged end users of the modelling tool voiced concerns regarding the scope and utility of introducing new software tools given their workload levels and the number of tools and new initiatives for other purposes demanding their attention. Any decisions regarding eventual tool implementation and deployment would have to

include extensive further engagement with those end users and provisions for soliciting and incorporating their feedback.

A key requirement for appropriate use of any modelling tool is knowledge or data concerning the resources (such as staff time, clinical environments and equipment, supplies and utilities) used in delivering different services. Readily available patient activity and administrative national data sources could be used to derive estimates of historical demand for services and the quantity of resources required to deliver these services to a patient population. Indeed, a calibration strategy was devised as part of this project using data from the Hospital Episode Statistics for patient activity and the NHS Information Centre for Health and Social Care workforce database for staffing. Inevitably, a number of parameters that need to be estimated are not available directly from data sources. Although, any existing data could potentially be augmented with the tacit knowledge of clinical experts and health managers, with recourse to guidelines and other clinical texts where appropriate, it is important to note that questions remain as to whether such a tool could be calibrated at a sufficiently detailed level for it to be useful. Maintaining and updating the calibration data would pose additional and substantial burdens, especially considering the dynamic and rapidly changing nature of the NHS.

A number of questions arise in the discourse about the role and scope of OR in projects with a strategic focus, as opposed to an operational one, and where the constituent elements of the problem are intrinsically unstructured. For example, can a simplified model ever capture enough of the essential details of the problem to meaningfully address it? On the other hand, is it practical or even possible to build a model that includes all necessary details and relaxes most of the restrictive modelling assumptions? Which exactly are the trade-offs between how comprehensive the model is, the level of detail within the model and the feasibility of model calibration?

In this paper, we have given a detailed account of an OR project that investigated the feasibility of developing a modelling tool that can be used to provide estimates of unmet demand for services under conditions of disruption while allowing for different levels of system response. The intended usage of such tool would be to provide emergency planning personnel with a systematic means of informing improvements to

NHS resilience and in assessing the potential impact on resilience of any proposed reconfiguration of services. Despite this important role, we identified a number of significant barriers to successful implementation and deployment including those arising from computational, calibration, maintenance, user acceptability and practicality issues.

A sensible way forward for a national or regional health care service interested in resilience planning would include an investment into creating and maintaining a data repository that will include detailed and current information on health care facilities (hospitals etc.), the local availability of resources (staff, supplies etc.) and the availability of resources and services that are outsourced to external organisations. Such a data repository, in addition to assisting in the operational decision during an emergency, should provide the foundation and data needed to then develop a modelling tool to assist in resilience planning.

Our work demonstrates the intrinsic difficulty, and yet potential value, of using OR to support policy-makers in addressing complex strategic questions for which there is no obvious, well-structured method or approach. Informing such debates demands a more nuanced approach than that of modelling to support a well-defined operational decision. There is a lack of explicit, clear thinking about resilience and the strategic consequences of service reconfiguration at present: the work discussed in this paper was used to good effect in identifying and presenting different elements of this problem both to policy colleagues and regional planners and exploring the potential viability of developing and putting to use a software tool. In this regard, our work made a beneficial contribution to the decision making process, which, ultimately concluded that such a tool was infeasible at present without significant additional investment.

Acknowledgements

We gratefully acknowledge the insights and advice we received from the following individuals: Sally Bell, Tom Bowen, Mary Gallivan, Dr Chooi Lee, Guy Mayou, Sam Roberts, Prof Tom Treasure, Andy Wapling.

References

Aaby K, Herrmann J W, Jordan C S, Treadwell M, and Wood K (2006). Montgomery County's Public Health Service uses operations research to plan emergency mass dispensing and vaccination clinics. *Interfaces* **36**(6): 569-579.

Academy of Medical Royal Colleges (2007). *Acute health care services - Report of a working party*. Academy of Medical Royal Colleges: London.

Altay N and Green W G (2006). OR/MS research in disaster operations management. *European Journal of Operational Research* **175**(1): 475-493.

Brandeau M L, McCoy J H, Hupert N, Holty J E, and Bravata D M (2009). Recommendations for modeling disaster responses in public health and medicine: a position paper of the society for medical decision making. *Med Decis Making* **29**(4): 438-460.

Cabinet Office (2013). UK Resilience website: a resource for civil protection practitioners, supporting the work which goes on across the United Kingdom to improve emergency preparedness. <http://update.cabinetoffice.gov.uk/ukresilience>, accessed 5 March 2013.

Cabinet Office (2008). National Risk Register – 2008 edition: Updating the public on the government's current assessment of the potential impact of a range of different civil emergency risks. <https://www.gov.uk/government/publications/national-risk-register-of-civil-emergencies>, accessed 11 July 2013.

Department of Health (2008). The NHS Resilience and Business Continuity Management (BCM) Guidance: Interim National Guidelines for NHS organisations. http://webarchive.nationalarchives.gov.uk/+/www.dh.gov.uk/en/Consultations/Liveconsultations/DH_085422, accessed 5 March 2013.

Eichner M, Schwehm M, Duerr H P, and Brockmann S O (2007). The influenza pandemic preparedness planning tool Influsim. *BMC Infectious Diseases* **7**.

Fulop N, Walters R, Perri, Spurgeon P (2011). Implementing changes to hospital services: Factors influencing the process and 'results' of reconfiguration. *Health Policy*. doi:10.1016/j.healthpol.2011.05.015.

Green L V and Kolesar P J (2004). Improving emergency responsiveness with management science. *Management Science* **50**(8): 1001-1014.

Hupert N, Mushlin A L, and Callahan M A (2002). Modeling the public health response to bioterrorism: Using discrete event simulation to design antibiotic distribution centers. *Medical Decision Making* **22**(5): S17-S25.

Imison C (2011). Reconfiguring hospital services. *London: King's Fund*. <http://www.kingsfund.org.uk/publications/articles/briefing-reconfiguring-hospital-services>, accessed 30 April 2013.

Jotshi A, Gong Q, and Batta R (2009). Dispatching and routing of emergency vehicles in disaster mitigation using data fusion. *Socio-Economic Planning Sciences* **43**(1): 1-24.

King's Fund (2007). *The reconfiguration of hospital services in England - A King's Fund briefing*. King's Fund: London.

Larson R C (2002). Public sector operations research: A personal journey. *Operations Research* **50**(1): 135-145.

Larson R C (2004). O.R. models for homeland security. *OR/MS Today* **31**(5): 22-29.

Larson R C (2006). Decision models for emergency response planning. In: Kamien D (Ed). *The McGraw-Hill Handbook of Homeland Security*. McGraw-Hill: New York, pp. 911-930.

Lee E K, Chen C H, Pietz F, and Benecke B (2009). Modeling and optimizing the public-health infrastructure for emergency response. *Interfaces* **39**(5): 476-490.

Pollock A M, Dunnigan M G, Gaffney D, Price D, and Shaoul J (1999). The private finance initiative - Planning the "new" NHS: downsizing for the 21st century. *British Medical Journal* **319**(7203): 179-184.

Rolland E, Patterson R A, Ward K, and Dodin B (2010). Decision support for disaster management. *Oper Manag Res* **3**:68-79.

Simpson N C and Hancock P G (2009). Fifty years of operational research and emergency response. *Journal of the Operational Research Society* **60**: S126-S139.

Tufekci S and Wallace W A (1998). The emerging area of emergency management and engineering. *IEEE Transactions on Engineering Management* **45**(2): 103-105.

Warwickshire County Council (2011) Warwickshire's Community Risk Register for 2010-2011. <http://www.warwickshire.gov.uk/communityriskregister>, accessed 11 July 2013.

Wein L M, Craft D L, and Kaplan E H (2003). Emergency response to an anthrax attack. *Proceedings of the National Academy of Sciences of the United States of America* **100**(7): 4346-4351.

Wein L M, Liu Y F, and Motskin A (2009). Analyzing the homeland security of the US-Mexico border. *Risk Analysis* **29**(5): 699-713.

Wright P D, Liberatore M J, and Nydick R L (2006). A survey of operations research models and applications in homeland security. *Interfaces* **36**(6): 514-529.

Zaric G S, Bravata D M, Holty J E C, McDonald K M, Owens D K, and Brandeau M L (2008). Modeling the logistics of response to anthrax bioterrorism. *Medical Decision Making* **28**(3): 332-350.

Appendix

In this section we introduce the notation used to characterise the demand for and delivery of health care services within a local health system and its capability to adapt service activity in the face of disruption or changed patterns of demand. We begin by describing the main elements associated with service activity. This is followed by descriptions of the supply and demand relationships which link these elements and the various levels of service activity. We then set out the integer programming problem that represents the most sophisticated response to disruption considered and discuss the computational feasibility of solving such problems. Finally, we describe the three heuristic algorithms each reflecting less sophisticated but perhaps more plausible rules for resource allocation in the face of disruption.

MODEL NOTATION

The following definitions summarise the elements used in the mathematical model, where possible following mnemonic conventions:

Let H denote the number of hospitals within the health system, indexed h .

Let C be the number of distinct care services, indexed c .

Let R be the number of distinct types of resources that are considered, indexed r .

For $1 \leq c \leq C$, $1 \leq h \leq H$, let $d'_{c,h}$ denote the demand (in terms of the number of cases presenting for treatment) over the period of interest for the c -th care service at the h -th hospital and $d_{c,h}$ denote the demand over the equivalent period prior to disruption.

For $1 \leq c \leq C$, $1 \leq h \leq H$, let $v'_{c,h}$ denote the number of cases requiring the c -th care service that are treated at the h -th hospital over the period of interest and $v_{c,h}$ denote the number treated over the equivalent period prior to disruption.

For $1 \leq h \leq H$, $1 \leq r \leq R$, let $q'_{h,r,1}$ and $q_{h,r,1}$ be the quantities of resource r that are available from on-site resource reserves at the h -th hospital over the period of interest

and over an equivalent period prior to disruption respectively. Similarly, let $q'_{h,r,2}$ and $q_{h,r,2}$ be the quantities available via the supply chain.

For $1 \leq h \leq H$, $1 \leq r \leq R$, let $T'_{r,h}$ and $T_{r,h}$ denote the total amount of resource r available at the h -th hospital during the period of interest and during an equivalent period prior to disruption respectively.

For $1 \leq c \leq C$, $1 \leq r \leq R$, let $x_{c,r}$ denote the quantity of resource r required, per case, for the delivery of care service c .

For $1 \leq c \leq C$, $1 \leq h \leq H$, let $w_{c,h}$ denote a user-defined weight reflecting the perceived importance of providing the c -th care service at the h -th hospital.

RESILIENCE PLANNING BASED ON UNMET DEMAND

The focus of our analysis is the unmet demand that occurs as a result of a disruption and the extent to which this can be mitigated by the choice of the number of services to provide in each hospital. In the model we express the unmet demand under disruption for care service c at the h -th hospital as $d'_{c,h} - v'_{c,h}$. The corresponding weighted unmet demand is $w_{c,h}(d'_{c,h} - v'_{c,h})$.

This unmet demand can be summed over all services and hospitals to give the aggregated weighted unmet demand as follows:

$$U = \sum_{h=1}^H \sum_{c=1}^C w_{c,h}(d'_{c,h} - v'_{c,h}). \quad (1)$$

It is important to note that, due to our assumption that demand is met prior to the period of disruption, any unmet demand can be attributed to the disruption.

As stated previously, in order to assess impact, a view has to be taken as to what might constitute a plausible response to disruption on the part of the local health system. We formulated a number of possible responses, chosen as they characterise

responses that have varying degrees of sophistication. The first formulation is based on an assumption that resources would be allocated to services within each hospital in order to minimise unmet demand. There follow three different heuristic algorithms characterising plausible rules for resource allocation in the face of disruption.

OPTIMISATION-BASED RESPONSE TO DISRUPTION

We begin this section by discussing a number of constraints that are necessary given the nature of the model.

We note that, in terms of the notation defined above, the following relationships follow from the definitions of the total amounts of different resource types available at each hospital over the period of interest:

$$T'_{r,h} = q'_{r,h,1} + q'_{r,h,2} \quad 1 \leq h \leq H, 1 \leq r \leq R. \quad (2)$$

Given that the number of cases treated over the period of interest cannot exceed the number of cases presenting for treatment and that both are non-negative, there are the following constraints:

$$0 \leq v'_{c,h} \leq d'_{c,h} \quad 1 \leq c \leq C, 1 \leq h \leq H. \quad (3)$$

Care activity is also constrained in terms of resource availability, summarised by the following:

$$\sum_{c=1}^C x_{c,r} v'_{c,h} \leq T'_{r,h} \quad 1 \leq h \leq H, 1 \leq r \leq R, \quad (4)$$

which, in view of (2), can be re-written as

$$\sum_{c=1}^C x_{c,r} v'_{c,h} \leq q'_{r,h,1} + q'_{r,h,2}, \quad 1 \leq h \leq H, 1 \leq r \leq R. \quad (5)$$

An integer programming formulation has been used to express the problem of minimising unmet weighted demand by reallocating resources and determining levels of different forms of care service within each hospital, by minimising the function U , given in (1), subject to (3) and (5) and the requirement that the decision variables $\{v'_{c,h}\}$ are non-negative integers.

We note that although such modelling implicitly allows for reallocation of resources within a single hospital, there are no explicit decision variables related to such reallocation, although the numerical values associated with the optimal allocation can be inferred from the solution to the integer programming problem. We also note that it is theoretically straightforward to extend this approach to include the case where resources or demand can be reallocated between hospitals, although at the expense of a considerably heavier computational burden (as seen in the results of an investigation in the next sub-section). In reality such a response would be challenging given that it would require coordinated decision making across hospitals and thus, have opted for the more simple case of within-hospital reallocation of resources and demand.

RESULTS OF COMPUTATIONAL FEASIBILITY TESTS

Key to assessing the impact of disruption is to gauge what might constitute a plausible response to disruption on the part of the NHS. The mathematical framework for the impact assessment tool thus introduced various approaches for assessing impact given different levels of response. There are limits to the ability of a software tool to solve the proposed optimisation problems associated with minimising unmet demand. These limits depend on the number of resources, assets, supply chain assets and services that are included in the model, as well as hardware and software limitations.

To assess the computational practicalities of the proposed approaches, we compared the number of constraints and decision variables for various optimisation problems associated with the measurement of disruption and feasible response under illustrative but realistic assumptions about the numbers of hospitals, supply chain assets, services and resources (Table A1). The first row of the table refers to optimisation within individual hospitals, with no reallocation of resources or demand between them; the resulting problem is one that could be solved using widely available software. The

next three rows refer to optimisation problems concerning the reallocation of resources, demand or both across a cluster of five local hospitals; the resulting problems could potentially be solved but perhaps only by using specialist software. The final three rows refer to the reallocation of resources, demand or both across an entire region; the resulting problems would test or exceed the capabilities of specialist software.

Table A1: Size of the optimisation-based disruption measure problems.

	No of hospitals	No of supply chain assets	No of services	No of resources	No of decision variables	No of constraints
Within-hospital optimisation	1	N/A	15	60	15	105
Resource reallocation	5	10	15	60	5027	3375
Demand reallocation	5	10	15	60	601	150
Resource and demand reallocation	5	10	15	60	5103	3450
Resource reallocation	50	10	15	60	39452	33750
Demand reallocation	50	10	15	60	6001	1500
Resource and demand reallocation	50	10	15	60	40203	34500

HEURISTIC RESPONSES TO DISRUPTION

We describe below three simple heuristics devised to characterise differing levels of potential response. Once more, we assume fixed resource and demand allocation at each hospital.

Algorithm 1 – Single step scaling of service activity

This algorithm represents a ‘business-as-usual’ approach until the first resource constraint is reached. It calculates the care activity under disruption by scaling down

the pre-disruption activity by the minimum proportion that is feasible (given the disrupted resource availability). Expressed mathematically, the measure of disruption

is obtained by defining $\rho_h = \min_r \frac{T'_{r,h}}{T_{r,h}}$ and then setting $v'_{c,h} = \min(\rho_h v_{c,h}, d'_{c,h})$ for

$1 \leq c \leq C$. This algorithm is relatively simple and easy to implement, although no reallocation of resources between services is incorporated within this response and so it is likely to over-estimate impact. We have opted to include it in the analysis and the prototype modelling tool as it provides a computationally efficient method of obtaining an initial lower bound to the problem.

Algorithm 2 – Iterative scaling of service activity (resource based)

The main idea behind this algorithm is to consider each resource in turn and scale down the caseload delivery by an appropriate factor to reflect the reduction in resource availability. In what follows, let K be the set of all resources and L be the set of care services. For each care service c , define L_r as the subset of L consisting of all c such that $x_{c,r} > 0$ (i.e. the set of all care services that depend on resource r).

The algorithm is carried out on a hospital-by-hospital basis, so the full algorithm consists of carrying out the steps set out below for each hospital (beginning, for each hospital, with K and L as the full set of resources and services respectively).

Step 1: From the set K , pick r_{\min} such that $\frac{T'_{r,h}}{T_{r,h}}$ is minimised. Calculate the

scaled care activity for all $c \in L_{r_{\min}}$ at h , conditional on the reduced resource

availability, by setting $v'_{c,h} = \min\left(\left[\frac{T'_{r_{\min},h}}{T_{r_{\min},h}} v_{c,h}\right], d'_{c,h}\right)$ for $1 \leq c \leq C$.

Step 2: Update resource availability assuming the delivery described in Step 1

takes place. Thus, for each resource r , replace $T'_{r,h}$ by $T'_{r,h} - \sum_{c \in L_{r_{\min}}} v'_{c,h} x_{c,r}$.

Step 3: Remove the resource r_{\min} from the set K and remove all care services c such that $x_{c,r_{\min}} > 0$ from L . If either K or L is empty, the algorithm terminates.

Otherwise, return to Step 1.

Algorithm 3 – Iterative scaling of service activity (service-based)

The motivation for this algorithm is an attempt to prioritise care services according to the user-defined weights $w_{c,h}$) as far as possible by, potentially, redirecting resources away from services deemed to be of lower priority. The algorithm is carried out on a hospital-by-hospital basis, so the full algorithm consists of carrying out the steps set out below at each hospital (i.e. for each value of h).

Step 1: From the set of weights $\{w_{c,h} \mid 1 \leq c \leq C\}$ associated with h , choose the largest member and calculate the maximum feasible activity of the associated service c at h , which is given by $\min_{r \in K_c} \left\lfloor \frac{T'_{r,h}}{x_{c,r}} \right\rfloor$ (where $K_c = \{r \mid x_{c,r} > 0\}$ i.e. the set of resources r required to deliver care service c). The expression for the service delivery is then given by $v'_{c,h} = \min \left(\min_{r \in K_c} \left\lfloor \frac{T'_{r,h}}{x_{c,r}} \right\rfloor, d'_{c,h} \right)$, which precludes the possibility of service over-provision. The notation $\lfloor z \rfloor$ refers to the integer part of z .

Step 2: Update all resource availability assuming the delivery described in Step 1 takes place. Under this scenario, the resource availability is reduced for all resources required in the delivery of c . This means that we replace $T'_{r,h}$ by

$$T'_{r,h} - v'_{c,h} x_{c,r} \text{ for each resource } r \in K_c.$$

Step 3: Remove $w_{c,h}$ from the set of weights associated with hospital h . If the set of weights is empty then terminate the algorithm. Otherwise, return to Step 1.

This algorithm implicitly assumes that all weights are distinct. In the case of equal weights, the steps are carried out for all different permutations of the services carrying equal weights and the permutation which corresponds to the minimal unmet demand is chosen.

SUPPLEMENT

In this supplement, we summarise the numerical inputs for the illustrative example described in the main paper and the assumptions and data sources used to calculate these inputs. The main data sources used for parameter estimation, where possible, were the NHS Information Centre for Health and Social Care workforce database (<http://www.ic.nhs.uk/>), used to gauge the size of the health service workforce, and the Hospital Episode Statistics (HES) database (<http://www.hesonline.nhs.uk/>), used to estimate the volume of cases by medical specialty. We also used the published findings of previous studies as well as expert opinion. We use a ‘finished consultant episode (FCE)’, which is defined by HES as a period of care under the responsibility of a consultant of a given medical specialty, as a proxy for a case of a given care service.

CARE SERVICE DEFINITIONS

The service groupings used in the example correspond, as far as possible, to the definitions used by the NHS Information Centre for Health and Social Care workforce database, with, in some cases, adjustments to the definitions based on differences in categorisation between the HES data and the workforce database. The breakdowns are as follows:

- Paediatric services consist of Paediatrics and Paediatric cardiology
- Maternity services consist of Obstetrics and Gynaecology
- Surgery consists of Cardiothoracic surgery, General surgery, Neurosurgery, Paediatric surgery, Plastic surgery, Trauma and orthopaedic surgery, Urology, Ophthalmology and Otolaryngology (the last two placed in the surgical group according to the Workforce Database).
- Acute medicine, for our purposes, comprises the following specialties: Ear, Nose and Throat, Accident and Emergency, General Medicine, Gastroenterology, Critical Care Medicine, Clinical Haematology, Clinical Genetics, Palliative Medicine, Respiratory (Thoracic) Medicine, Infectious Diseases, Nephrology, Medical Oncology, Neurology, Rheumatology and Geriatric Medicine. This categorisation is based on the General Medicine

group definition from the NHS Information Centre for Health and Social Care workforce database and information from the HES database.

RESOURCE REQUIREMENTS OF SERVICES

The numerical values for the resource requirement (per case of a given service) reported here are, in some sense, averages of the constituent specialties of each of the four aggregated service groups we consider. The medical staff figures for the resource requirements are based on the ratio between medical staff full time equivalents (FTEs) in a given specialty (obtained from the NHS Information Centre for Health and Social Care workforce database) and the number of FCEs for that specialty (obtained from the HES database), Table S1. These ratios are summarised in Table S2 and are taken to reflect the ‘medical staff intensity’ of each specialty group.

Table S1: Resource requirements (units of resource required to deliver a case of service)

Services	Resources	Resource requirements	
		Emergency	Non-emergency*
Acute medical	Bed-days	4.2 [¶]	2.1
	Nursing staff-days	6.5 [†]	3.3
	Paediatrician-days	-	-
	Physician-days	1.6	0.8
	Surgeon-days	-	-
	Obstetrician-days	-	-
	Units of red blood cells	0.1 [‡]	0.06
	Clean instruments (relative measure)	1	0.5
Maternity	Bed-days	1.2	0.7
	Nursing staff-days	3.6	1.8
	Paediatrician-days	-	-
	Physician-days	-	-
	Surgeon-days	-	-
	Obstetrician-days	0.8	0.4
	Units of red blood cells	0.05	0.03
	Clean instruments (relative measure)	1	0.5

Paediatric	Bed-days	2.8	1.4
	Nursing staff-days	4.2	2.1
	Paediatrician-days	1.9	0.9
	Physician-days	-	-
	Surgeon-days	-	-
	Obstetrician-days	-	-
	Units of red blood cells	0.1	0.06
	Clean instruments (relative measure)	1	0.5
Surgery	Bed-days	4.1	2.1
	Nursing staff-days	4.5	2.2
	Paediatrician-days	-	-
	Physician-days	-	-
	Surgeon-days	1.8	0.9
	Obstetrician-days	-	-
	Units of red blood cells	0.2	0.1
	Clean instruments (relative measure)	2	1

* All non-emergency care estimates were calculated based on the assumption that the resource usage for non-emergency care is half that for emergency care.

¶ Estimates of bed-days and clean instruments were based on expert opinion.

† All staff-related estimates were based on a combination of data obtained from the NHS Information Centre for Health and Social Care workforce database and HES with assumptions based on expert opinion.

‡ Estimates of blood usage were based on a combination of data obtained from HES, NHS Blood and Transplant and a published paper (Wells et al., 2002), with assumptions based on expert opinion.

Table S2: Staff intensity per specialty group

Specialty	Relative number of medical staff in each specialty (adjusted so that it is equal to 1 for surgery)
Maternity	0.5
Paediatric	1.0
Surgery	1.0
Acute	
Medical	0.9

Nurse FTEs are defined relative to medical staff FTEs, assuming that each case (FCE) requires a contribution from both Medical staff and Nursing & Midwife staff, Table S3.

Table S3: Relative nursing to medical staff FTEs

Specialty	Ratio of nursing & midwife staff FTEs to medical staff FTEs
Maternity	4
Paediatric	2
Surgery	2
Acute Medical	4

The resource requirement for blood was obtained by combining NHS Blood and Transplant (NHSBT) data with published data (Forster *et al.*, 2003) outlining red blood cell unit usage by different specialties. Table S4 summarises the blood resource requirement. For simplicity, we have chosen to use the usage of red blood cell products as a proxy for all blood use.

Table S4: Blood resource requirement

Specialty	Units of blood per case
Surgery	0.2
Maternity	0.05
Paediatric	0.12
Acute medical	0.12

DEMAND

Table S5 summarises approximate relative demand for the different services, expressed as a percentage of the overall demand for all services, based on nationwide annual HES figures for FCEs. Also listed is the percentage of that demand that is attributed to emergency care (again, based on nationwide annual HES figures).

It is assumed that these relative proportions are constant at all locations and that the number and relative size of hospitals within the local health system are as in Table S6.

Table S5: Breakdown of admissions by specialty

Service	% of total demand	% of which is attributable to emergency care
Acute	48	49
Maternity	15	65
Paediatrics	10	28
Surgery	27	37

Table S6: Hospitals within the putative local health system

Hospital Category	Small	Medium	Large
Relative size (total admissions)	1	1.5	2
Number of hospitals modelled	3	5	2

These assumptions allow us to further break down the headline FCE figures on a per-hospital and per-service basis. These disaggregated demand figures were used as the demand input for the model runs.

RESOURCE AVAILABILITY

In order to populate the pre-disruption resource availability at each hospital in the illustrative example, we assumed that it was equal to the pre-disruption demand for resources (calculated from the resource requirements per case and the case demand inputs described above), with 5% excess capacity.

REFERENCES

Wells A W, Mounter P J, Chapman C E, Stainsby D, Wallis J P (2002). Where does blood go? Prospective observational study of red cell transfusion in north England. *British Medical Journal* **325**: 803.

Forster A J, Stiell I, Wells G, Lee A J, van Walraven C (2003). The effect of hospital occupancy on emergency department length of stay and patient disposition. *Academic Emergency Medicine* **10**: 127-133.