Using System Dynamics to Plan Investment in Alcohol Services

Douglas McKelvie, Eric Wolstenholme, Steve Arnold, David Monk

The Symmetric Partnership

The Grain Store, 127 Gloucester Road, Brighton, United Kingdom BN1 4AF

Tel: (0044) 131 557 0559, (0044) 1273 811 092

douglas.mckelvie@symmetricsd.co.uk, eric.wolstenholme@symmetricsd.co.uk,
steve.arnold@symmetricsd.co.uk, david.monk@symmetricsd.co.uk

Abstract

The authors developed a Dynamic Whole System Model of Alcohol Harm Reduction for the England Department of Health, to support local commissioning of alcohol related services.

The project used group model building, based on the best evidence available. It is intended to help local health commissioners reduce hospital admissions attributable to alcohol. The main "high impact" interventions incorporated are 'brief advice' in Primary Care, the employment of Alcohol Health Workers in hospital, and Specialised Treatment. The key output measures are hospital admissions and costs.

The model uses four consumption groups (Abstainers; Lower Risk; Increasing Risk; Higher Risk) including 'binge' and 'dependent' drinkers in more than one state. Each state has a differing propensity for hospital admissions. The model provides a dynamic cost analysis; as interventions move people between states hence changing their risk of admission to hospital. The model contains a set of policies parameterised by the Department of Health, but also allows for local settings.

The work relates to the search for consistent and cohesive policies by which central government can guide local actions. The approach of using dynamic models goes beyond action lists for guidance and allows localities to learn what will work for them.

Key Words: alcohol consumption, system dynamics, health, modelling, UK

Context and Client

The England Department of Health's Alcohol Improvement Programme (AIP) had a emit to promote a range of interventions designed to tackle alcohol related harm caused by excessive alcohol consumption. Tackling excessive consumption is a worthwhile end in itself; additionally, the AIP wanted to test the economic case for such investment, whether savings made as a result of health improvement, particularly reductions in admissions to hospital and attendances at accident and emergency departments, might be greater than the cost of the new services.

Research into these interventions, mainly in the form of follow-up studies of those treated, suggested that savings were possible. Informed by this evidence, the AIP developed in-house a spreadsheet model, known as the "Ready Reckoner", as part of a suite of tools to be used by local service commissioners (typically at local health organisation level). Typical behaviour of this model showed a linear relationship between projected savings and investment in new services.

In 2009, the AIP engaged a consultancy to develop a system dynamics tool that was built on a less simplified set of processes and assumptions, could be better adapted to fit local circumstances and requirements; and could be operated over a longer time frame. The project proceeded using an expert-group model building approach. Once the model was developed and further tested, it was developed into a runtime version designed for local use.

This modelling project continues a growing trend in health and social care analysis using dynamic simulation methods (Homer (2004), Hirsch et al (2009), Milstein et al (2010), Dangerfield (1999), Wolstenholme (1993, 1996, 2005) and Wolstenholme et al (2007, 2008, 2009, 2010)), particularly in the area of public health and condition management and treatment.
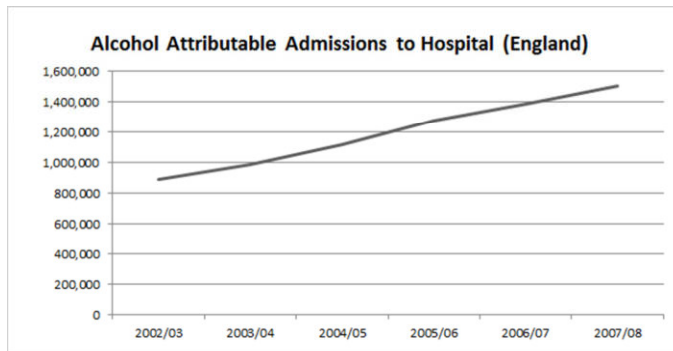
The Problem

Drug, alcohol and addiction problems have been studied using system dynamics. The primary concerns of these models have included the supply chain, the issue of addiction, problems caused by anti-social behaviour resulting from intoxication, and even modelling the biology by which alcohol is absorbed into the blood stream (Moxnes and Jensen, 2009).

The particular issue exercising this client was more mainstream, the wide range of damage done to the health of the general population through excessive consumption. Most, not all, of the interventions modelled here focus on reducing consumption, as a much wider concern than tackling dependent drinking (addiction) or episodic binge drinking. The Department of Health has sponsored considerable research into how this excessive consumption feeds its way into hospital admission. More than 40 types of hospital admission are found to be partly or wholly attributable to alcohol consumption.

The overall position is shown in Figure 1, showing the rise in "alcohol attributable" hospital admissions. This rise is attributed to increasing consumption combined with a rising adult population.

Figure 1 Alcohol Attributable Admissions to Hospital Per Annum

One way of tackling this problem is the provision of brief advice, combined with brief treatments (such as a series of four sessions with an adviser), to people believed to be drinking excessively. These services might be targeted on people attending accident and emergency departments, people recently admitted to hospital, or through some basic screening of patients consulting general practitioners in primary care.

There is a very large body of research evidence supporting such opportunistic case finding for alcohol misuse and the delivery of simple advice including at least 56 controlled trials (Moyer et al., 2002). A recent Cochrane Collaboration review (Kaner et al., 2007) provides substantial evidence for the effectiveness of alcohol case identification and the delivery of brief advice. For every eight people who receive simple alcohol advice, one will reduce their drinking to within low-risk levels (Moyer et al., 2002). This compares favourably with smoking where only one in twenty will act on the advice given (Silagy & Stead, 2003).

Evidence suggests that those drinking at increasing or higher risk who receive brief advice are twice as likely to moderate their drinking 6 to 12 months after brief advice when compared to drinkers receiving no brief advice (Wilk et al, 1997). Brief advice can reduce weekly drinking by between 13% and 34%, resulting in 2.9 to 8.7 fewer mean drinks per week with a significant effect on movement towards lower-risk alcohol use (Whitlock et al, 2004).

The dynamic hypothesis is that tackling consumption through these brief interventions will reduce alcohol consumption in the population, leading to health improvements that will show themselves in a reduction in alcohol-attributable hospital admissions, such that, in time, investment in new services is recovered from subsequent savings.

In the AIP's pre-existing spreadsheet tool, that causal chain is encapsulated in a simple statistical relationship between an input (brief interventions) and an output (reduction in hospital admissions). The effect is shown in the screenshot in Figure 2 where some

interventions (recruitment of alcohol health nurses) produce a linear reduction in hospital admissions. A key feature of this behaviour is that when a constant change is introduced, there is an immediate stepped reduction (actual against trend) followed by a rise at the same gradient as the projection.

Figure 2 Part of the Spreadsheet "Ready Reckoner"
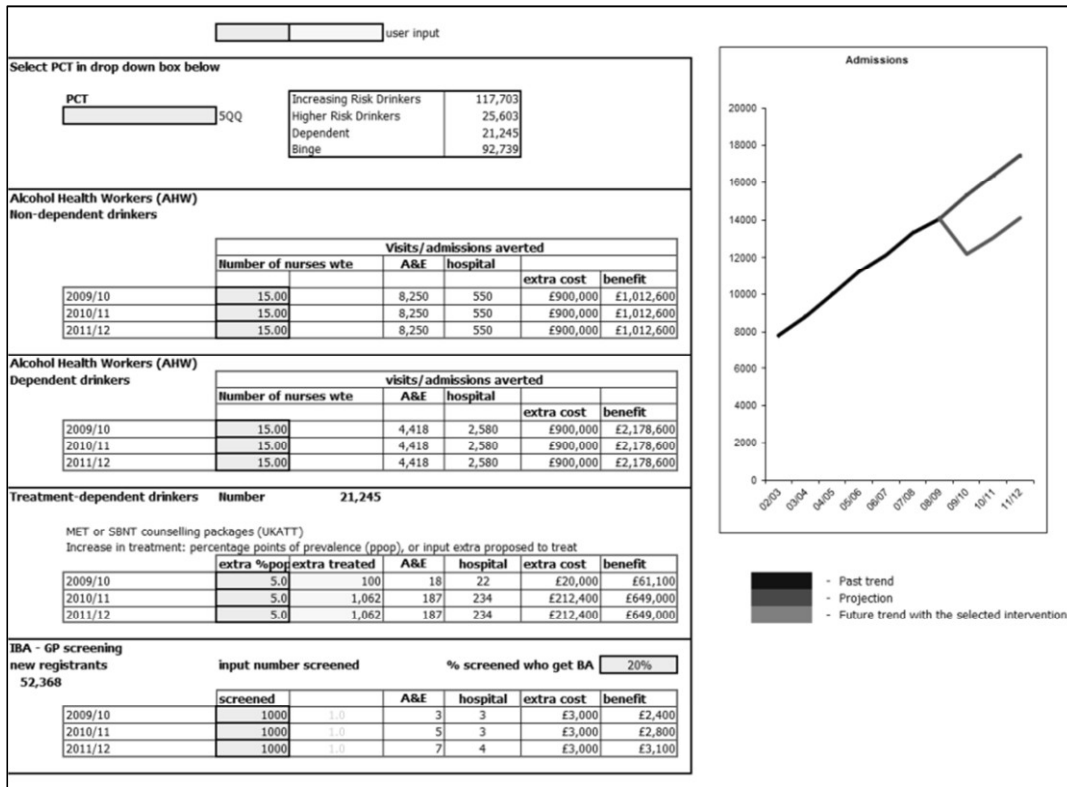
Method

The method adopted to build the model was group model building. The Department of Health assembled an expert group consisting of experts from the fields of public health, alcohol / substance abuse, health economics, drawn from academia, service commissioning, practice, policy development, operations research / data analysis and the "not for profit" sector. The group met at monthly intervals for eight months, facilitated by two consultants, both with a background of working in the health and care sectors combined with expertise in building system dynamics models. The model was developed using iThink software. Although there were a few changes in membership, the group composition remained reasonably consistent; this contributed greatly to the success of the project.

After the model was completed (and no model is ever finished) a smaller, core group continued to meet, testing the model under a range of scenarios and, ultimately, refining some of its structure, assumptions and functionality.

During this phase, the model was converted into a "runtime" version, with detailed documentation and online help menus, and tested with a wide range of possible users of the model, mainly local service commissioners. The model interface was adapted in the light of their feedback. At the time of writing, the model is due to be made widely available to service commissioners throughout England.

Description of Model

The main stock-flow structure of the model is the adult population of a locality; in the development phase, we used a notional population of 250,000 (rising to just under 300,000 in 20 years). The four stocks represent the consumption groups (abstinent, lower risk, increasing risk and higher risk), defined according to average weekly consumption of alcohol in units. This classification is widely used by the Department of Health for policy planning purposes. Information for the groups is available from the General Lifestyle Survey[1], an annual survey of a sample of households in Great Britain and is derived from questions about amounts of different drinks usually drunk on any one day during the last 12 months. The classification differs from international classifications of hazardous and harmful drinking based on the Alcohol Use Disorders Identification Test (AUDIT).

The benefit of using this classification is that much government planning data draws on it, but it does differ from international classifications of hazardous drinking.

Although the focus on consumption is foremost, the framework does contain within it other policy issues, such as dependence / addiction and binge drinking. A snapshot of how the adult population of England is distributed across these groups is shown in Figure 3.

Figure 3 Percentage in Each Consumption Group, also Showing How Dependent and Binge Drinkers are Distributed across Consumption Groups


---
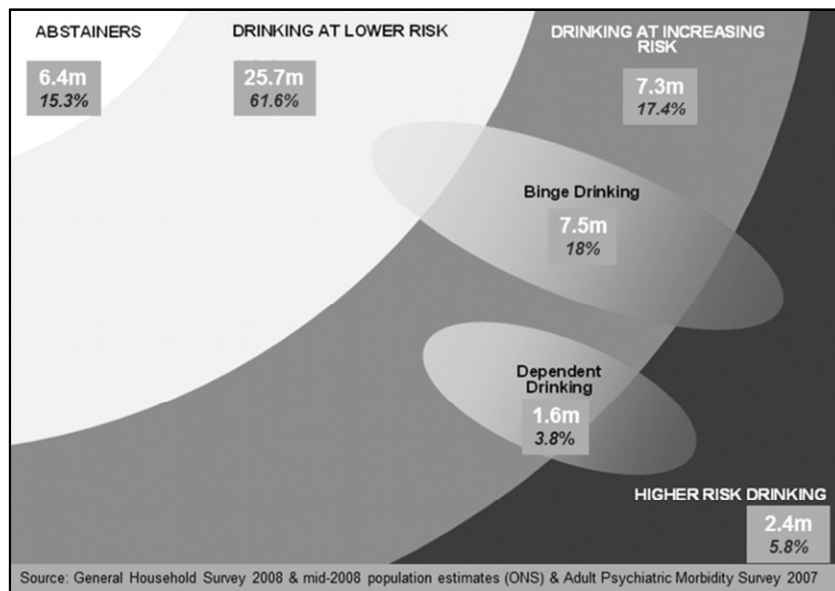
[1] Formerly known as General Household Survey, as in Figure 3

The main stock-flow chain represents people moving between these consumption groups. It bears some resemblance to system dynamics models of chronic health problems (Homer et al, 2004), with the crucial differences that:-

- People can move up <u>and</u> down the chain; in most chronic conditions models only degeneration is possible
- Movement is not just to the adjacent stock; people can move between any stocks (for example movement from the "higher risk" group to "abstinent" would be a requirement for someone following a "12 steps" approach)

Figure 4 Main Stock Flow Chain for Alcohol Consumption
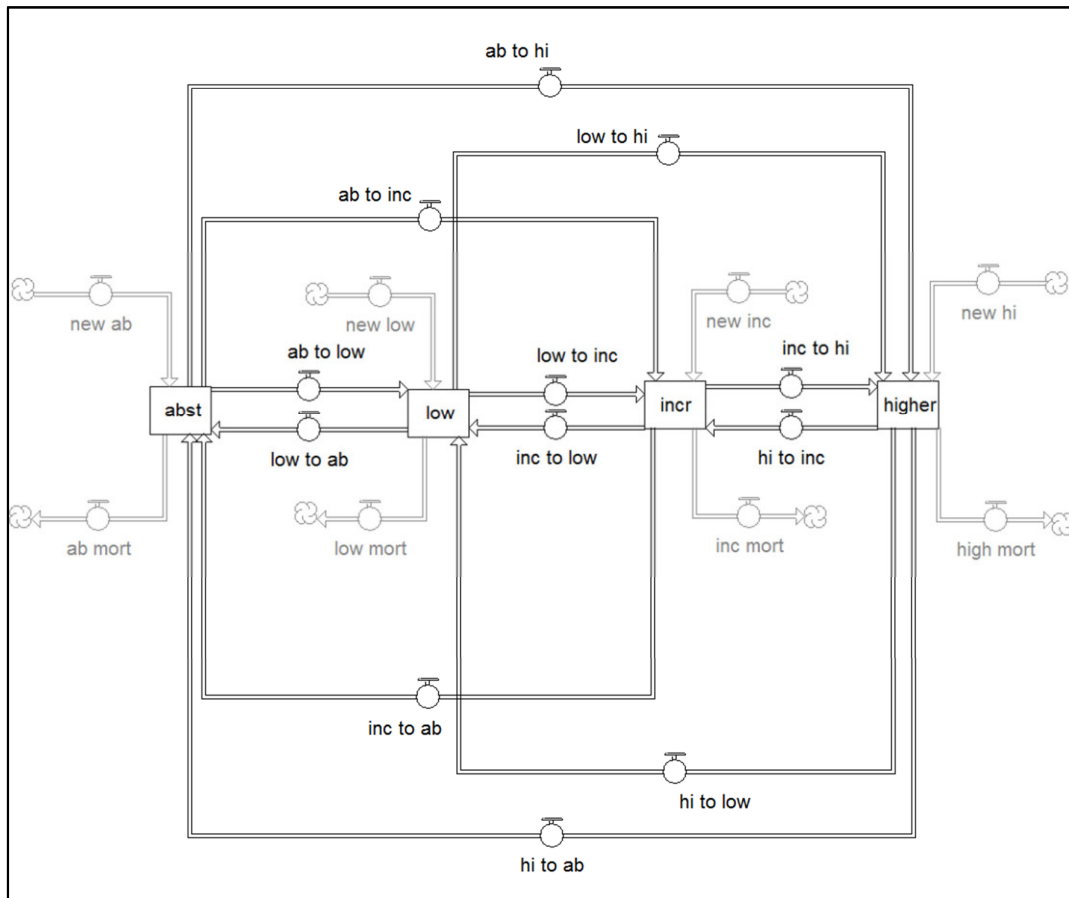
It is important to establish the main dynamics of this chain before applying the impact of additional interventions. These include:-

- The model represents the entire adult (18 and over) population of a discrete area
- Therefore, at any point in time everyone is in one stock; the states are mutually exclusive and exhaustive
- The total population is given as an exogenous input, and for most areas is projected to increase over time

- The model runs for 20 years with a time step of days; it is now clear that months would have been a satisfactory time step but in the initial development period some stakeholders were interested in some detailed analysis of hospital episodes, whose lengths of stay are best measured in days (as with most models built in a consulting environment, this one contains more detail than is necessary, and time does not allow for subsequent refinement and simplification)
- National data are available to give the initial distribution between consumption groups
- Adjustments in population are based on different mortality rate inputs for each group, combined with a stock adjustment of "new" cases (mainly people reaching the age of 18)
- To make that adjustment, we need to know the distribution by consumption group of people at age 18 (people enter the model into any of the consumption groups)
- Further adjustments to population are made based on the effectiveness of treatment interventions (the population might therefore grow to a larger number than the exogenous input for scenarios showing particularly effective interventions because more people are in the lower consumption groups which have lower mortality)
- A key requirement is to estimate the "base" rate of movement between the stocks (see below)

Estimates of the historical distribution of the population between the four stocks are available at national level, and the Department of Health had made its own projections of the future distribution (not using system dynamics). There had been a steady rise in alcohol consumption in the 1990s which reached a plateau around 2004; this plateau was projected to continue[2].

The client was able to undertake a secondary analysis of the General Household Survey data, comparing people's consumption in two successive years (this was not straightforward because the data set is of households rather than individuals). This estimated the annual rate of movement (in terms of the percentage of people who moved from each stock to each other stock in the course of a year) as follows:-

---

[2] As things have turned out, there is some evidence of a reduction in consumption (perhaps as a consequence of the economic downturn)

Figure 5 Percentage of People who have Moved from Each Consumption Group to Each Other Group after One Year

| From \ To	Abstinent	Lower Risk	Increasing Risk	Higher Risk
Abstinent		25	0	0
Lower Risk	5		8	0
Increasing Risk	1	21		11
Higher Risk	1	10	19	

The model is designed to be operated by a variety of users, and it is possible to select from a range of starting years. Typically, the model has been run using 2005 as the starting point, but to produce a reference mode, the model was run from 2000. Without any additional dynamics resulting from the introduction of service interventions, the model's base-case behaviour is determined by:-

- The flow rates set out in Figure 5
- The initial percentage in each stock
- Assumptions about the distribution of people joining the model (drinking status on 18[th] birthday [3])
- Different mortality rates applied to each stock
- Population growth

These inputs combined to produce a good match between the percentage distribution of people between drinking groups generated by the model and the (reference) historical plus projected data held by the Department of Health. Figure 6 compares the model outputs for the percentage in each stock in the base run against the reference data.

---

[3] Obviously, not based on behaviour on the day itself, the legal drinking age in England

Figure 6 Reference Mode Comparing Base Run of Model Output for Percentage in Each Consumption Group against Reference Data
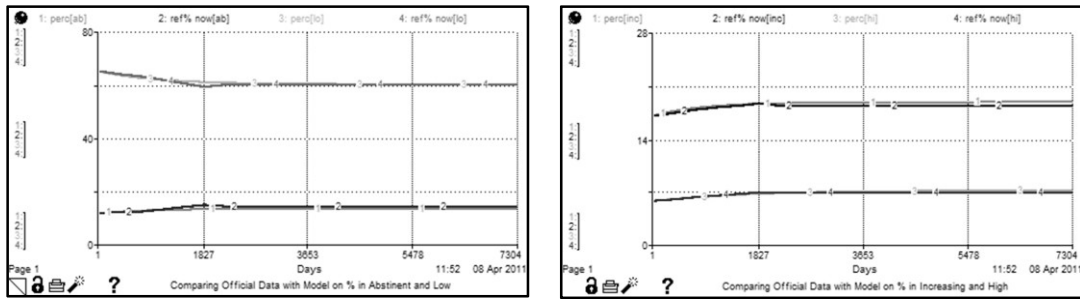
The total adult population is growing, as are alcohol attributable hospital admissions:-

Figure 7 Total Adult Population and Alcohol Attributable Admissions in Model Base Run
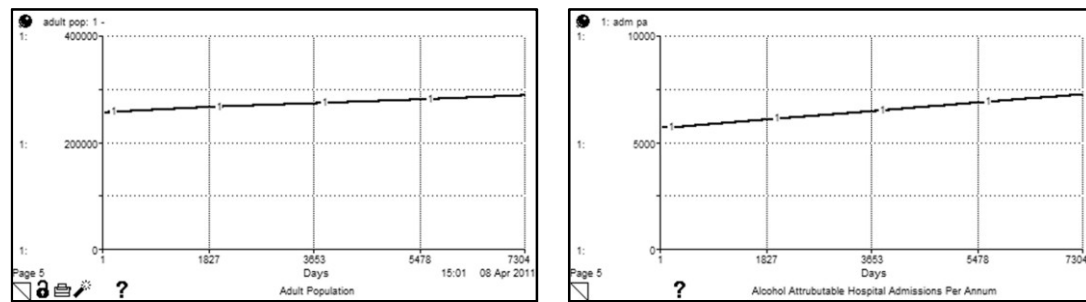
The two main output variables of concern are:-

- Alcohol attributable hospital admissions
- Cost of services (including hospital and interventions, described below)

The purpose of the model, therefore, is to explore the extent to which spending money on specific interventions generates reductions in hospital admissions and therefore saves money. The causal mechanisms are that:-

- Alcohol attributable hospital admissions vary by consumption group and are higher in the higher consumption groups
- Service interventions will increase the rate at which people move into lower consumption groups
- This takes place against a background of considerable churn between consumption groups in the base case
- With more people in the lower consumption groups, alcohol admissions will also reduce

9

This is summarised below.

Figure 8 Diagram Showing Assumptions behind Causal Relationships in the Model
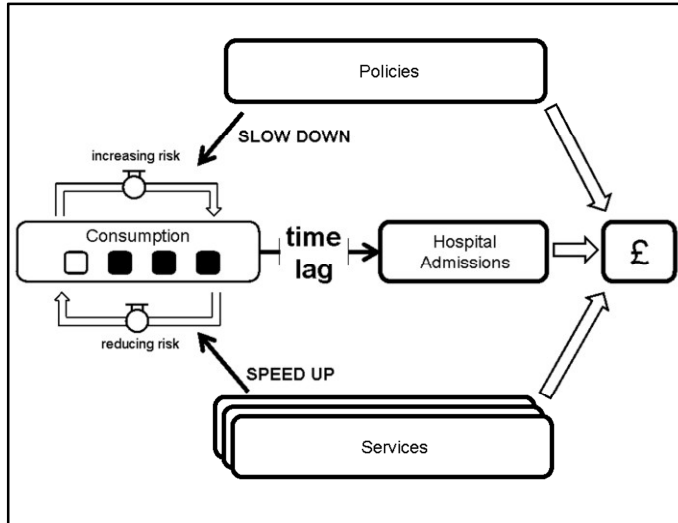
This diagram is neither a stock-flow map nor a causal loop diagram. In the Consumption part, people move between consumption groups. The distribution between groups determines the rate of hospital admissions, but there is also a delay here, because reductions in consumption are not immediately followed by improvements in health.

By commissioning interventions (services) it is possible to change the flow rates in the main model. After the time delay, this will show itself as a reduction in admissions.

There is a further possibility, capable of being represented in the model but politically highly sensitive, that "policies" could be introduced to encourage people to drink less. In the United Kingdom, that debate has centred on the feasibility of introducing minimum unit pricing. This model was developed during the lead-in time to a general election, when there was no political will to explore such options in detail. The matter remains very much up for debate.

(Note that although the diagram makes it look like services exclusively operate by speeding up the reducing consumption flows, and policies by slowing down the rate of increase, actually each will do a bit of both).

Hospital admissions are a mixture of acute (e.g. accident and injury) and chronic conditions (related to, or exacerbated by alcohol consumption). It is possible to set a different time lag for each category of admission[4]. For example, following a reduction in consumption someone's propensity to have an accident is reduced almost immediately, whereas chronic liver disease will continue (though some recovery is possible).

---

[4] Here, 44 types of admission have been collapsed into 10 categories in the model, using arrays

The financial implications are that both hospital admissions and services cost money. "Policies" may not cost a great deal, or may cost the Department of Health nothing (being enforced by other means). Increased expenditure on services is incurred immediately, whereas any reductions in expenditure resulting from reductions in hospital admission will take time to show (because of the delay factor).

There are several service interventions modelled. They all have a similar structure:-

- A demand function based on the main stock – flow chain or some other part of the model
- A capacity (the size of the investment, how much of the intervention to procure)
- An effect (the percentage of people receiving the intervention who reduce their consumption as a result) which feeds back into the main stock – flow chain
- A cost (based on the capacity utilisation)
- A setting determining how far into the model run the service should be switched on (typically the model is run for five years, then the intervention is introduced)

The characteristics of the main interventions are:-

"Identification and brief advice" is provided in primary care. People visiting a general practitioner for any reason are asked some routine screening questions. Depending on the response to these questions, some supplementary advice about alcohol consumption is offered.

Brief interventions provided by "Alcohol Health Workers (AHWs) in hospital" (Owens, 2005 and 2007), in which people admitted to hospital, where there is reason to belief this might be alcohol-related, are offered brief counselling from a specialist (who might be a nurse). In this case, the maximum possible users of the service is governed by the alcohol related hospital admissions figure that is itself generated by the model. Most model users to date have allocated a relatively modest capacity to this service, reflecting current reality that such a service tends to be targeted on dependent drinkers.

"Specialised services" – treatments for dependence, have tended to be less well used in modelling sessions. They are the more traditional services geared towards tackling dependence / addiction, characterised by being typically more expensive and provided in smaller numbers. As the modelling work progressed, so the focus in the model was more on the "brief intervention" services which are tackling consumption.

Typical Model Behaviour

In the base case of the model, the overall population is rising, and the distribution of people between the consumption groups has plateaued as described above. The only expenditure being incurred is on hospital admissions, and also some specialised services.

The model dashboard enables a user to navigate to different service settings, and also to control which services are to be used in a particular model run through the use of switches.

Figure 9 The Model Dashboard
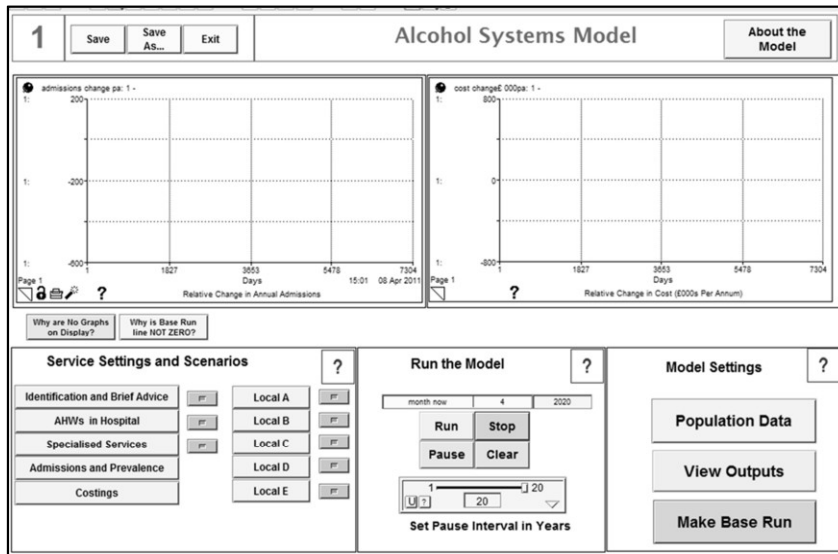
The Service Settings consist of buttons enabling a user to navigate to a series of input screens for each type of service. As well as the three services identified above, the buttons headed Local A – E enable a user to describe up to five more locally-defined services (in terms of capacity, cost, effect on consumption). The "admissions and prevalence" settings do not have a switch because hospital services cannot be anything other than "on".

The Model Settings lead to:

- a series of population input screens that are mostly set-up once (to match the area being modelled) but not changed once scenarios are being run
- a series of output screens on which different graphs are displayed
- a "make base run" function (explained below)

1 Base Run
In the sequence of runs that follows, the start year is 2005 and the model runs for 20 years.

The base run of the model is normally a run with no additional services switched on.

At this point, we are most interested in the outputs for:-

- number of hospital admissions per annum
- total cost per annum

In the base case, both variables are rising (with similar gradients, because hospital costs are the main costs).

**Figure 10 Total Admissions and Costs Outputs from Base Run**
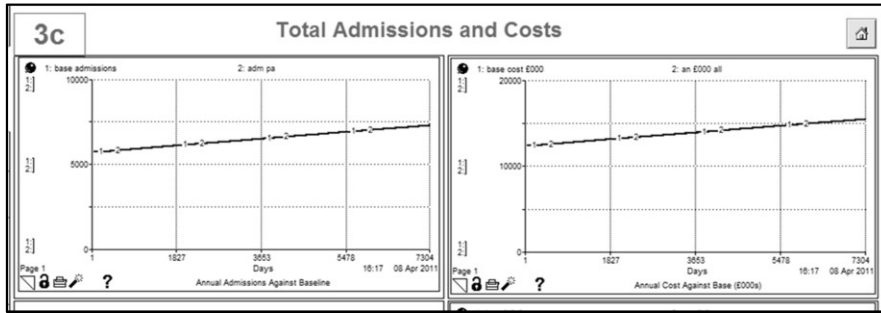
These graphs show the current run (admissions per annum or expenditure per annum in £000s) mapped against variables called "base". Since this is the base run these lines are identical. Both admissions and costs are rising, mainly because of a rise in population (consumption having reached a plateau).

2 Scenario using Identification and Brief Advice in Primary Care
In this scenario, a new intervention is switched on five years into the run. It consists of:-

- 10% of primary care practices screening all "new registrations" where 8% of the population registers (normally following a change of address) with a practice per annum
- 5% of primary care practices screening every patient attending where 65% of the population visit a practice at least once per annum (people are screened only once in the year)
- Those identified as being in the two higher consumption groups are offered advice as a result of which 12.5% reduce their consumption enough to move into a lower group

**Figure 11 Admissions and Costs Comparing Scenario 2 against Base**
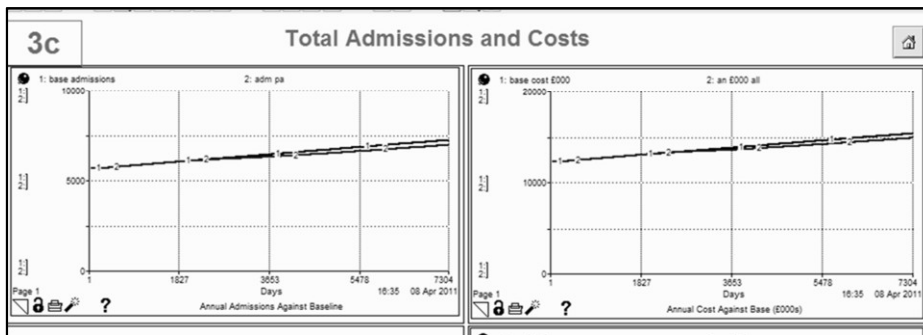
At this point, it becomes quite hard to discern much change. The intervention point is five years into the model (day 1827 on the graph), and after a considerable time, some reduction in admissions, then costs, begins to show.

Faced with these outputs, model users were finding it difficult to evaluate what the model was telling them. And in the model's earlier days these graphs were

"comparative" graphs (plotting one variable over a series of runs) which tended to get quite cluttered.

What they really wanted was a graph of the gap between the current run and the base run, an output not readily generated in iThink. By exporting several key variables (such as admissions and costs) to a spreadsheet, then importing these values into graphical variables called "base variable…" this problem was overcome, and new sets of variables plotting the difference between the current run and base were generated. These variables were ultimately selected to be the main variables displayed, as comparative graphs on the dashboard:-

Figure 12 Model Dashboard showing "Difference between this Run and Base Run" Graphs for Admissions and Costs (Base and Scenario 2)
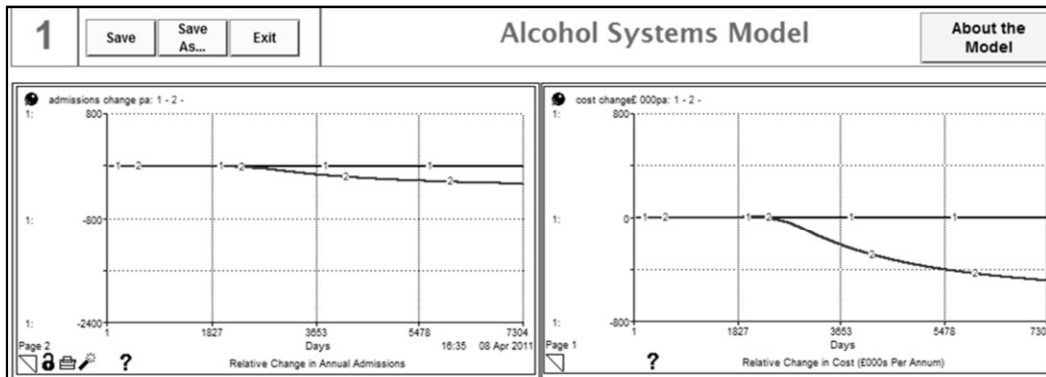

Now, the base run is run 1, and the scenario is run 2. Obviously, in the base run, the difference between the variable and the base is 0. In the scenario, it becomes clearer that there is a reduction in admissions which is quite evident about two years after the introduction of the intervention, and that costs increase slightly (because when the intervention is first introduced there is a cost but no saving yet on admissions). These patterns become more obvious in the next scenario.

3 Scenario: Introduce Alcohol Health Workers in Hospital
In this scenario, four workers are employed to provide brief advice and interventions, typically a series of four appointments spread over 10 weeks, to people attending hospital (either accident and emergency / emergency room, or admitted) who are believed to have an alcohol attributable admission. As a result of the intervention, 60% of people seen reduce their consumption and move into lower consumption groups.

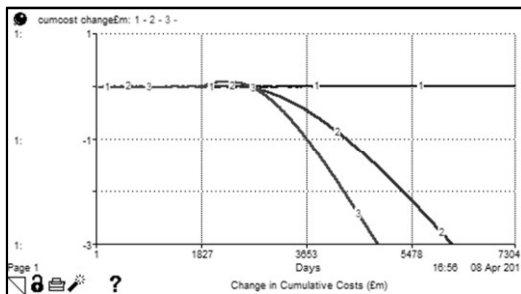**Figure 13 Difference in Admissions and Costs for Base, Run 2 and Run 3**
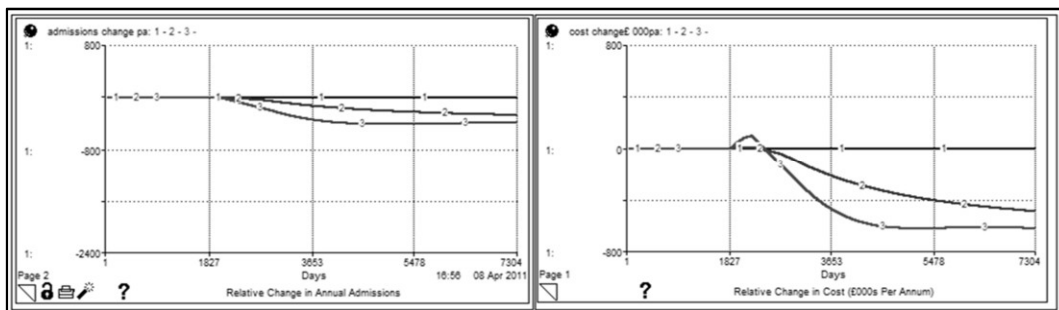
Most of the outputs produced by the model take this shape, an initial increase in expenditure (because a new service is introduced), followed by a subsequent reduction in expenditure as the reduction in admissions produces cost savings that are greater than the new service expenditure.

For the most part, service commissioners have been content to compare expenditure levels year on year, and are content to consider the point at which the change in expenditure line crosses the X-axis as a "breakeven" point.

Strictly speaking, this is not the case, and some commissioners suggested that the breakeven point is the point at which cumulative change in cost is negative (so far, nobody has asked for a discounted cash-flow calculation).

The cumulative cost change graph (generated in the same way, by the export / import of base cumulative costs through a spreadsheet) for these model runs is harder to scale but shows the breakeven point for scenario 3 to be within the third year following the change.

Figure 14 Difference in Cumulative Cost Compared with Cumulative Cost of Base Run for the same runs

4 Scenario: Combining Both Interventions
Obviously, it is possible to switch on more than one service intervention. Run 4 represents the combined effect of using the services introduced in Runs 2 and 3.

Figure 15 Difference in Admissions and Cost Graphs for All Runs
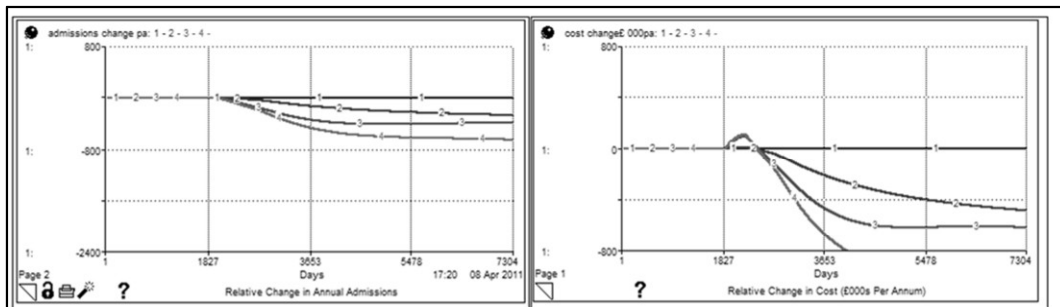
Making Sense of the Findings

From the client's point of view, the model (disappointingly) showed less of a reduction in admissions (hence fewer savings) than the existing spreadsheet, which functioned essentially by extrapolating from a statistical relationship between service interventions and subsequent admission rates drawn from various studies. Because of this, the model was subjected to a prolonged period of testing and questioning, which established that the model faithfully represented the collective mental model of the client. The systems model could more readily represent the causal connections, delays and accumulations that had hitherto merely been captured within a statistical relationship, in particular:-

- The basic assumption that alcohol related ill health is caused by consumption – therefore those in the higher consumption groups will tend to experience more hospital admissions than those in the lower group
- That specific interventions will move people down the consumption chain
- Nevertheless, there is a time delay between change in consumption and improvement in some chronic conditions
- There is a cumulative effect because interventions move people into a different stock (thereafter, some will move back up, or continue down, the consumption chain)

There are other connections, for example, the supply of interventions located in hospital can be varied (e.g. by recruiting more staff) but demand (maximum number of alcohol attributable admissions) is a function of model behaviour (we might run out of people to treat). In the spreadsheet it is theoretically possible to go on recruiting staff, with a linear reduction in admissions, to the point where future admissions go negative because the number of admissions "saved" exceeds the actual number of admissions (an impossibility). In the systems model, the obvious "limit to growth" in the demand function means that beyond a certain amount, there is no point in expanding a service.

Ironically, because few will have tested the spreadsheet under these extreme conditions, the systems model tended to show that a much greater investment in hospital-based Alcohol Health Workers was indicated than any of the users who developed and piloted the model envisaged.

In the course of a long dialogue between the client and the modelling team seeking to explain the modest performance of the systems model, two further refinements were made:-

- It would be possible that some people might reduce their consumption within a group but not enough to move into a lower stock
- When modelling the impact of interventions on people admitted to hospital with chronic conditions, using a treatment effect on the general population (only a fraction of whom have such a condition) dilutes the impact

Further adjustments were made using co-flows, including a representation of prevalence and incidence of the chronic conditions modelled, which, as expected, improved the overall treatment effect in the systems model, but still showing fewer savings than were projected by the spreadsheet.

Wider Applications

Costs and Benefits
The model has been used so far to compare the impact of policy interventions on "costs". The model does not, as yet, make any attempt to quantify "benefits", although these benefits can be shown.

The most obvious benefit is that the general health of the population would improve if alcohol consumption goes down. In the last run of the model shown, combining brief interventions in hospital and in primary care, in addition to the main outputs shown (admissions and costs), it is worth looking at the distribution of the population by consumption group.

Figure 16 Number in Higher and Increasing Risk Consumption Groups Compared against Base Run
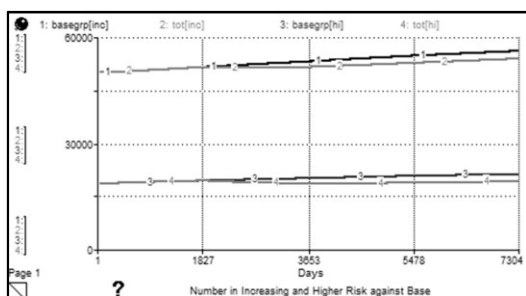
In this model run, out of a final population of 297,000, the number in the Increasing Risk group has reduced from a base 56,500 to 54,300, and the number in the Higher Risk group from 21,200 to 19,100, a total reduction of over 4,000 from the two top consumption groups.

The model does not at present quantify these population health improvements, such as in Quality Adjusted Life Years.

Another obvious benefit would look beyond the health sector. Such a reduction in consumption would also be shown in the criminal justice system, given a link between certain types of crime and alcohol consumption. There is interest in exploring this further but it is beyond the scope of the present enquiry.

Non-Treatment Approaches to Tackling Alcohol Consumption
The model is currently limited to exploring the impact of specific treatment interventions on alcohol consumption. The model could also be expanded to include interventions not currently modelled, for example the impact of different pricing interventions on health and criminal justice outcomes.

The model does not offer a properly calibrated representation of the sensitivity of alcohol consumption to pricing, but it can be assumed that the intention of any minimum-pricing policy would be, in stock / flow terms, to reduce the rate at which people flow up the consumption chain, and increase the rate at which people flow down the chain.

Introducing a very simple reduction to the flow rates, of "assume that the increase-in-consumption flows reduce by 10% and the reduction-in-consumption flows increase by

10%", after five years, produces a dramatic effect on hospital admissions and hence costs. Figure 17 shows these outputs.

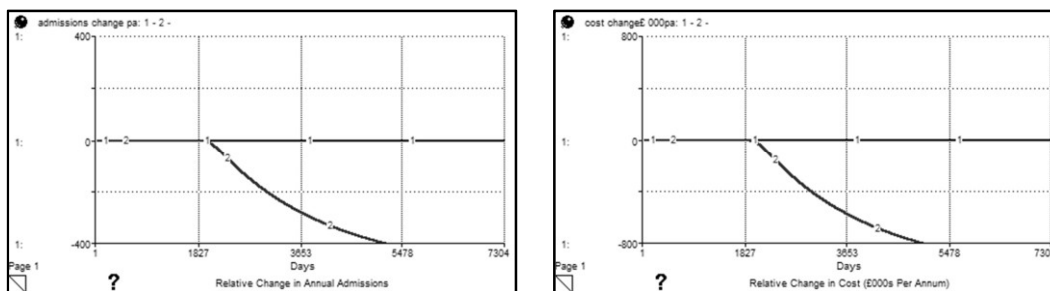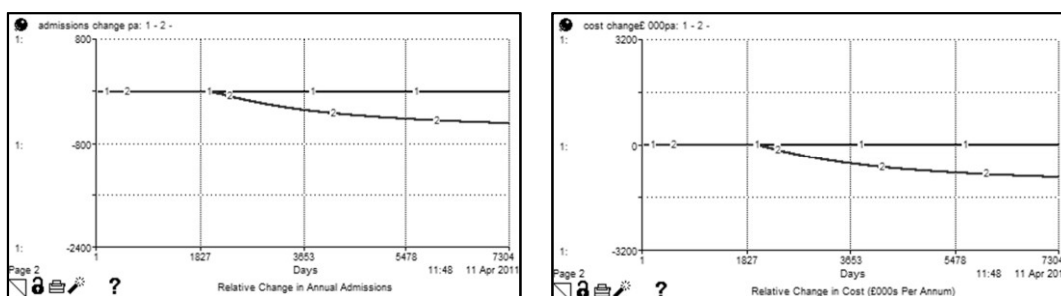**Figure 17 Effect of "Minimum Pricing"**

Figure 18 The same graphs on a different scale

The most striking difference between these outputs and the outputs generated by the treatment intervention scenarios is that there is no rise in cost following the intervention point, because the policy would be introduced at no cost to the health system. There is no "investment" in services to be made.

Making the Model Usable

The model has required an amount of secondary analysis by the client in order to estimate inputs for such variables as the "normal rate of movement between consumption groups". In addition, although a great deal of detailed analysis of alcohol attributable hospital admissions exists, largely based on the development of "alcohol attributable fractions" by which (all) admissions data are multiplied to report the number of alcohol attributable admissions by area, the model generated new data requirements, such as the need to have these admissions broken down by consumption group.

Although it would undoubtedly be possible to construct a simpler, higher level model, showing only one category of admissions and only one treatment intervention (applied in varying degrees), to create a meaningful tool for service commissioners it is necessary to disaggregate admissions by type (by using an array function), and to separate out different types of intervention (partly because they each have different

18

demand functions, such as people attending primary care, people admitted to hospital, people who are alcohol-dependent).

The model is also linked to various spreadsheets, both to create more comprehensive data reports and also to achieve the "create a base run" routine (involving exporting and importing a data set).

The model is now available as an "isee Runtime" version, which has been tested by a range of potential users. The runtime model interface has a range of features designed to make the model usable by service commissioners, including help menus, options to enter local data, the possibility of choosing from a variety of possible start years.

Conclusion

The model development has proceeded from an initial requirement for a tool to be used to communicate the benefits of a range of service interventions, so that local service commissioners might make more informed choices about future investment.

In the course of that development, which was supported by national experts in the field, questions arose about why the model projections differed somewhat from those that had been expected. This provoked a lengthy period of questioning of, learning from, and improvement of, the model.

The outcome is a tool that is capable of providing a more sophisticated analysis of the likely cost implications of new service investment, in a financial climate where increased investment in any area of public health may need to be justified by a business case showing overall savings. The model shows that any investment will be followed by a period in which more money is spent on services; in time the whole system of treatment might cost less, and it is possible to compute both the time at which expenditure will be less than baseline, and also when the initial investment might be recovered.

The systems model arguably provides a more compelling picture of these phenomena than a spreadsheet approach that extrapolates from a particular statistical relationship, because it represents the subtleties of the impact of delays between changes in consumption and changes in health, as well as some limits to growth.

The model is capable of informing a wider debate, for example about the implications for the criminal justice system of these health interventions. There would be merit in extending the range of policy options modelled to include alcohol-pricing; a system dynamics approach would enhance that debate.

Finally, the model interface has been designed to make it usable by a wide range of (non-expert) users. The question of how much support such users will need and how it might be afforded given the current state of the public finances remains open, and the authors would hope to return to some of these themes of making complex models more accessible at a future date.

References

Dangerfield B.C., System Dynamics Applications to European Health Care Issues", Journal of the Operational Research Society, Vol. 50, No. 4, 1999. pp345-353.

Hirsch G, Homer J. System Dynamics Applications to Health Care in the United States. In *Encyclopaedia of Complexity and System Science*, ed. R Meyers, Springer-Verlag, Berlin, Germany. ISBN 978-0-387-75888-6. July 2009.

Homer J., A System Dynamics Model for Cocaine Prevalence Estimation and Trend Projection *Journal of Drug Issues*, 23(2): 251-279, 1993.

Homer J. Projecting the Impact of Law Enforcement on Cocaine Prevalence, *Journal of Drug Issues*, 23(2): 281-295, 1993.

Homer J. A System Dynamics Model of National Cocaine Prevalence, *System Dynamics Review*, 9(1): 49-78, 1993.

Homer J. A Dynamic Model of Cocaine Prevalence in the United States, In *System Dynamics*, ed. Y Barlas, *Encyclopaedia of Life Support Systems* (EOLSS), available at http://www.eolss.net. Developed under auspices of UNESCO, EOLSS Publishers, Oxford, UK. 2004.

Homer J, Hirsch G, Minniti m and Pierson M (2004) Models for Collaboration; how System Dynamics Helped a Community Organize Cost-effective Care for Chronic Illness, *System Dynamics Review,* 3, 20

Kaner E, Beyer F, Dickinson H, Pienaar E, Campbell F, Schlesinger C, Heather N, Saunders J, Bernand B. Brief interventions for excessive drinkers in primary health care settings. Cochrane Database of Systematic Reviews 2007, Issue 2. Art No.: CD004148 DOI: 10.1002 /14651858.CD004148.pub3.

Milstein B, Homer J, Hirsch G. Analyzing National Health Reform Strategies with a Dynamic Simulation Model, *American Journal of Public Health*, 100(5):811-819, 2010.

Moxnes E, Jensen L 2009 Drunker than Intended: Misperceptions and Information Treatments, *Drug Alcohol Dependence* 2009 Nov 1 ; 105(1-2), pp 63-70

Owens, L (2005) The Burden of Alcohol on the NHS. Pharmacology & Therapeutics, The University of Liverpool. Doctor of Philosophy

Moriarty, K. Cassidy, P. Dalton, D. Farrell, M. Gilmore, I. Hawkey, C. Keaney, F. Moore, K. Owens, L. Rhodes, J. Shenker, D. Sheron, N. (2010) Alcohol-Related disease; Meeting the challenge of improved quality of care and better use of resources. A joint position paper of behalf of the BSG BASL and Alcohol health Alliance UK

Moyer, A., Finney, J., Swearingen, C. and Vergun, P. (2002) Brief Interventions for alcohol problems: a meta-analytic review of controlled investigations in treatment - seeking and non-treatment seeking populations, Addiction, 97, 279-292.

Owens, L. Morton, D. (2007) Effective nurse-led and patient-centred alcohol interventions. Nursing in Practice July/August: 50-55.

Owens, L. (2005) How do NHS Hospitals in England deal with patients with alcohol-related problems? A questionnaire survey. Alcohol & Alcoholism. 40 (5): 409-12.

Silagy, C. and Stead, L.F. (2003) Physician advice for smoking cessation (Cochrane Review), in: The Cochrane Library, Issue 4 (Chichester, Wiley).

Whitlock, E.P., Polen, M.R., Green, C.A., Orleans, T. and Klein, J. (2004), Behavioralcounseling interventions in primary care to reduce risky/harmful alcohol use by adults: a summary of the evidence for the US Preventive Services Task Force. Annals of Internal Medicine, 140, 557-568.

Wilk, A.I., Jensen, N.M. and Havighurst, T.C. (1997), Meta-analysis of randomized control trials addressing brief interventions in heavy alcohol drinkers, Journal of General Internal Medicine, 12, 274-283.

Wolstenholme, E.F., A Case Study in Community Care using Systems Thinking, *Journal of the Operational Research Society*, Vol. 44, No. 9, September 1993, pp 925-934.

Wolstenholme, E.F., A Management Flight Simulator for Community Care, in Enhancing Decision Making in the NHS, Ed. S. Cropper, 1996, Open University Press, Milton Keynes.

Wolstenholme, E.F., A Patient Flow Perspective of UK Health Services, *System Dynamics Review.* Vol. 15, no. 3, 253-273.(1996)

Wolstenholme, E.F., The potential of system dynamics – a new era for strategic planning. The NHS Confederation Future Healthcare Network, Leading Edge Series, Issue 10 November 2005.

Wolstenholme, E.F., Monk D, McKelvie D and Smith G "Influencing and Interpreting Health and Social Care Policy in the UK" **,** *In Complex Decision Making: Theory and Practice,* Eds. Qudrat-Ullah H, Spector M and Davidsen P., New England Complex Systems Institute Book Series on Complexity, Springer-Verlag. (2007**)**

Wolstenholme, E.F., Arnold S Monk D Todd D and McKelvie D., Coping but not Coping in Health and Social Care - masking the reality of running organisations well beyond safe design capacity. *System Dynamics Review* 50[th] Anniversary Celebratory Edition Vol.23 No.4 ((2008)

Wolstenholme, E.F., System Dynamics Applications to Health Care in the United Kingdom and Europe in *Encyclopaedia of Complexity and System Science*, Springer Germany, pp 445-470. (2010)

Wolstenholme, E.F., Dynamic Cost Benefit Analysis for Mental Health Reform (with David Monk and David Todd) *Kybernetes,* Vol. 39 No. 9/10, 2010 pp. 1645-1658 Emerald Group Publishing Limited (2010)