Optimisation of system dynamics models

Dangerfield, BC

SYSTEM DYNAMICS MODELS, OPTIMISATION OF

Brian Dangerfield
Centre for OR & Applied Statistics
Salford Business School
Maxwell Building
University of Salford
M5 4WT
UK

(b.c.dangerfield@salford.ac.uk)

Article Outline

Glossary

Econometrics: a statistical approach to economic modelling in which all the parameters in the structural equations are estimated according to a 'best fit' to historical data.

Maximum likelihood: a statistical concept which underpins calibration optimisation and which generates the most likely parameter values; it is equivalent to the parameter set which minimises the chi-square value.

Objective function: see **Payoff** below

Optimisation: the process of improving a model's results in terms of either an aspect of its performance or by calibrating it to fit reported time series data.

Payoff: a formula which expresses the objective, say, maximisation of profits, minimisation of costs or minimisation of the differences between a model variable and historical data on that variable.

Zero-one parameter: a parameter which is used as a multiplier in a policy equation and serves the effect of bringing in or removing a particular influence in determining the optimal policy.

1. Definition of the Subject and Its Importance

The term ‘optimisation’ when related to system dynamics (SD) models has a special significance. It relates to the mechanism used to improve the model vis-à-vis a criterion. This collapses into two fundamentally different intentions. Firstly one may wish to improve the model in terms of its performance. For instance, it may be desired to minimise overall costs of inventory whilst still offering a satisfactory level of service to the downstream customer. So the criterion here is cost and this would be minimised after searching the parameter space related to service level. The direction of need may be reversed and maximisation may be desired as, for instance, if one had a model of a firm and wished to maximise profit subject to an acceptable level of payroll and advertising costs. Here the parameter space being explored would involve both payroll and advertising parameters. This type of optimisation might be described generically as *policy optimisation.*

Optimisation of performance is also the *raison d’etre* of other management science tools, most notably mathematical programming. But such tools are usually static: they offer the ‘optimum’ resource allocation given a set of constraints and a performance function to either maximise or minimise. These models normally relate to a single time point and may then need to be re-run on a weekly or monthly basis to determine a new optimal resource allocation. In addition, these models are often linear (certainly so in the case of linear programming) whereas SD models are usually non-linear. So the essential differences are that SD model optimisation for performance involves both a dynamic and a non-linear model.

A separate improvement to the model may be sought where it is required to fit the model to past time series data. Optimisation here involves minimising a statistical function which expresses how well the model fits a time-series of data pertaining to an important model variable. In other words a vector of parameters are explored with a view to determining the particular parameter combination which offers the best fit between the chosen important model variable and a past time series dataset of this variable. This type of optimisation might be generically termed *model calibration*. If *all* the parameters in the SD model are determined in this fashion then the process is equivalent to the technique of econometric modelling. A good comparison between system dynamics and econometric modelling can be found in Meadows and Robinson (1985).

2. Optimisation as calibration

In these circumstances we wish to determine optimal parameters, those which, following a search of the parameter space, offer the best fit of a particular model variable to a time series dataset on that variable taken from real world reporting.

As an example consider a variation of the one of the epidemic models which are made available with the Vensim™ software. The stock-flow diagram is presented as figure 1.

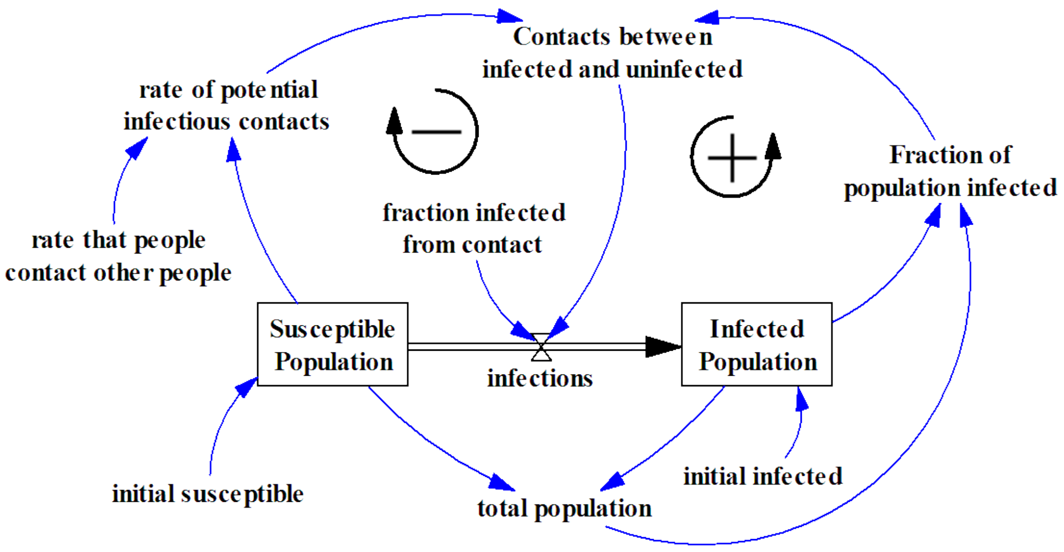

Figure 1: Stock-flow diagram for a simple epidemic model

In this epidemiological system members of a susceptible population become infected and join the infected population. Epidemiologists call this an S-I model. It is a simpler variation of the S-I-R model which includes recovered (R) individuals.

Suppose some data on new infections (at intervals of five days) is available covering 25 days of a real-world epidemic. The model is set with a time horizon of 50 days which is consistent with, say, a flu epidemic or an infectious outbreak of dysentery in a closed population such as a cruise ship. The 'current' run of the model is shown in figure 2, with the real-world data included for comparison.

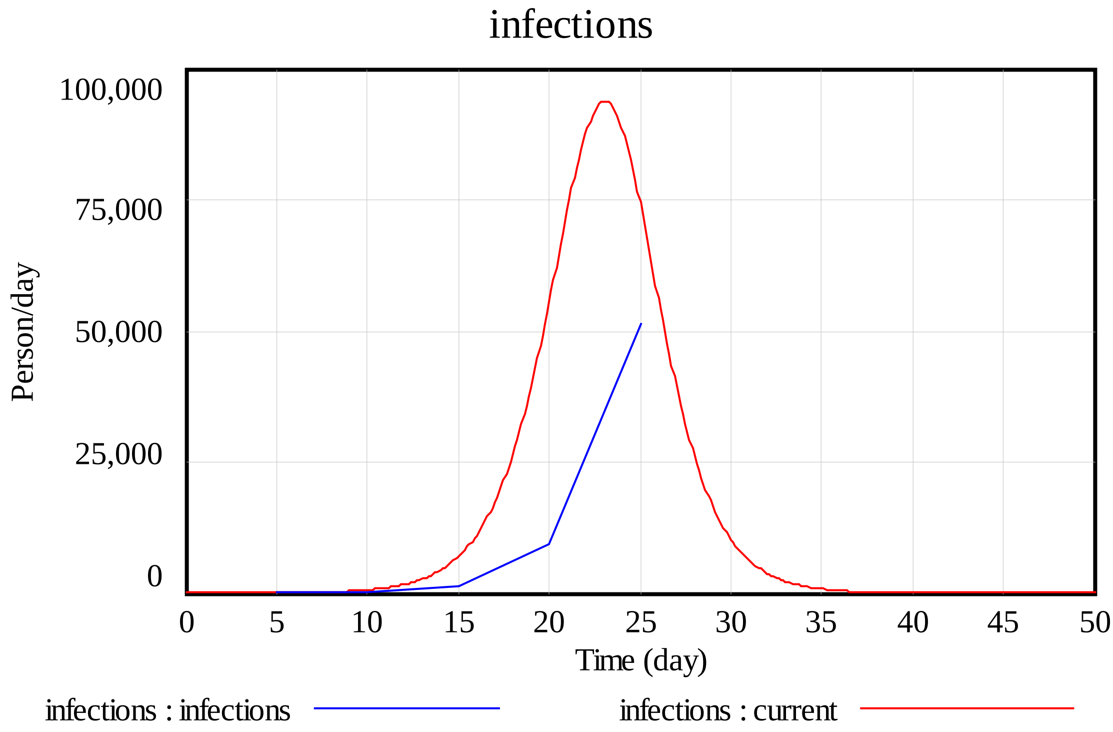

Figure 2: Current (Base) run of the model and reported data on infections

Clearly there is not a very good correspondence between the actual data and the model variable for the infection rate (infections). We wish to achieve a better calibration and so there is a need to select relevant parameters through which the calibration optimisation can be performed over. Referring back to figure 1 we can see that the *fraction infected from contact* and the *rate that people contact other people* are two possible parameters to consider. The initial infected and initial susceptible are also parameters of the model in the strict sense of the term, but we will ignore them on this occasion. In this model the *initial infected* is 10 persons and *initial susceptibles* number 750,000 persons.

The chosen value for the *fraction infected from contact* is 0.1 while that for the *rate that people contact other people* is 5.0. The former is a dimensionless number while the latter is measured as a fraction per day (1/day). This is obtained from consideration of the *rate of potential infectious contacts* (persons/day) as a proportion of the *susceptible population* (persons).

The optimisation process for calibration involves reading into the model the time series data, in this case on new infections, and, secondly, determining the range for the search in parameter space. There is usually some basic background knowledge which allows a sensible range to be entered. For instance, a probability can only be specified between 0 and 1.0. In this case we have chosen to specify the ranges as follows:

$$0.03 <= \text{fraction infected from contact} <= 0.7$$

2 <= rate that people contact other people <= 10

A word of warning is necessary in respect of optimising delay parameters. Because there is a risk of mathematical instability in the model if the value of DT (the TIME STEP) is too large relative to the smallest first-order delay constant, it is important to ensure the TIME STEP employed in the model is sufficiently small to cope with delay constant values which may be reached during the search of the delay parameter space. In other words ensure the minimum number for the search range on the delay parameter is at least double the value of the TIME STEP.

Maximum likelihood estimation and the payoff function

The optimisation process involves a determination of what are termed statistically as maximum likelihood estimates. In Vensim™ this is achieved by maximising a payoff function. Initially this is negative and the optimisation process should ensure this becomes less negative. An ideal payoff value, after optimisation, would be zero. A weighting is needed in the payoff function too but for calibration optimisation this is normally 1.0. Driving the payoff value to be larger by making it less negative has parallels with the operation with the Simplex algorithm common in linear programming. This algorithm was conceived initially for problems where the objective function was to be minimised. Its use on maximisation problems is achieved by minimising the negative of the objective function.

During the calibration search, Vensim™ takes the difference between the model variable and the data value, multiplies it by the weight, squares it and adds it to the error sum. This error sum is minimised. Usually data points will not exist at every time point in the model. Here the model TIME STEP is 0.125 (1/8th) but let us assume that reported data on new infections have been made available only at times t= 5,10,15,20 and 25 so the sum of squares operation is performed only at these five time points.

The data is shown as Table 1

Table 1: Data used for calibration experiment

Time	5	10	15	20	25
Infections	30	230	1400	9500	51400

The recording point for reported data

System dynamics models differentiate between stock and flow variables and the software used for simulating such models advances by a small constant TIME STEP (also known as DT). This has implications for the task of fitting real-world reported data to each type of system dynamics model variable. The issue is: at what point in a continuum of time steps should the reported data be recorded at? This is important because the reported data has to be read into the model to be compared with the simulated data. The answer will be different for stock and flow variables.

Where the reported data relate to a stock variable the appropriate time point for recording will be known. If it is recorded at the end of the day (say a closing bank balance) then the appropriate point for data entry in the model will be the beginning of the next day. Thus the first data point above is at time *t*= 5 (5.00) and would, if it were a stock, correspond to a record taken at the very end of time period 4.

However, if the data relate to a flow variable, as in the case of new infections here, the number is the total new infections which have occurred over the entire time unit (day, week, month etc.) and so there is a decision to be reached as to which time point the data are entered at. This is because the TIME STEP (DT) is hardly ever as large as the basic time unit which the model is calibrated in. The use of 5 (10, 15 etc) above implies that the data on new infections over the period of time *t*= 0 to *t*= 5 is compared with the corresponding model variable at time 5+1*DT (and the new infections over the period *t*=5 to *t*= 10 at time 10+1*DT etc). A more appropriate selection might be towards the end of the 5-day time period. Following the example above using a TIME STEP= 0.125, this might be at time 4+7*DT (that is at 4.875).

Calibration optimisation results

Based upon the data on new infections shown above and the chosen ranges for the parameter search the following output is obtained (Table 2) After 114 simulations the optimised values for our two parameters are shown to be 0.08 and 5.12 and the payoff is over 2500 times larger (less negative). Replacing the original parameters with the optimised values reveals the result shown in figure 3. To take things further we may wish to put confidence intervals on the estimated parameters. One way of accomplishing this is by profiling the likelihood and is described in Dangerfield and Roberts (1996a).

Table 2: Results from the calibration optimisation

Initial point of search.
  fraction infected from contact = 0.1.
  rate that people contact other people = 5.
Simulations = 1.
Pass = 0.
Payoff = -2.67655e+009.
---------------------------------.
Maximum payoff found at:.
  fraction infected from contact = 0.0794332.
 *rate that people contact other people = 5.11568.
Simulations = 114.
Pass = 6.
Payoff = -1.06161e+006.
---------------------------------.

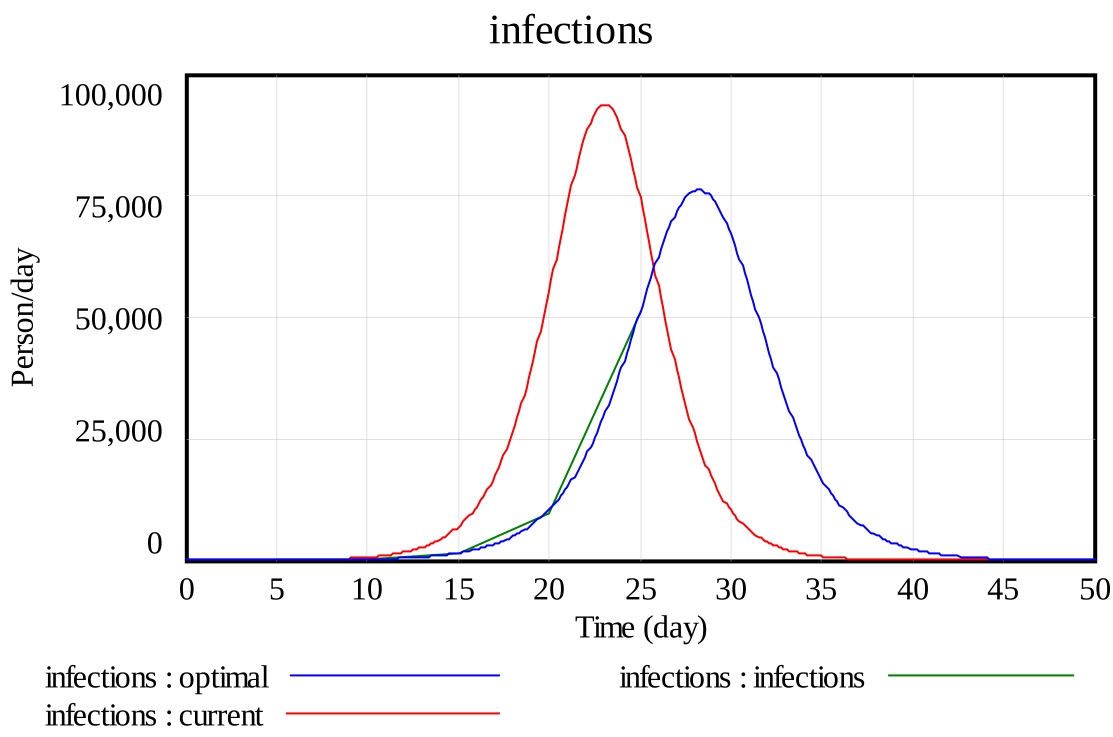

Figure 3: Reported data on infections and optimised (calibrated) model; the base case (current) is reproduced for reference

Avoid cumulated data

There might be a temptation to optimise parameters against cumulated data when the data is reported essentially as a flow, as is the case here. Were the data to be cumulated we would obtain as shown in Table 3.

Table 3: Cumulated reported data for the Infected Population

Time	5	10	15	20	25
Infected population	30	260	1660	11160	62560

The results from this optimisation are shown in Table 4. The ranges for the parameter space search are kept the same but the payoff function now involves a comparison of the model variable *infected population* with the corresponding cumulated data. Figure 4 shows the resultant fit to infected population is good but that is manifestly not borne out when we consider the plot of infections obtained from the same optimisation run (figure 5).

The reason for this is rooted in statistics. The maximum likelihood estimator is equivalent to the chi-squared statistic. This is turn assumes that each expected data value is independent. A cumulated data series would not exhibit this property of independence.

Table 4: Results from the calibration using cumulated data

Initial point of search
  fraction infected from contact = 0.1
  rate that people contact other people = 5
Simulations = 1
Pass = 0
Payoff = -2.48206e+011

Maximum payoff found at:
  fraction infected from contact = 0.0726811
 *rate that people contact other people = 4.96546
Simulations = 145
Pass = 6
Payoff = -212645

The final payoff is -212645

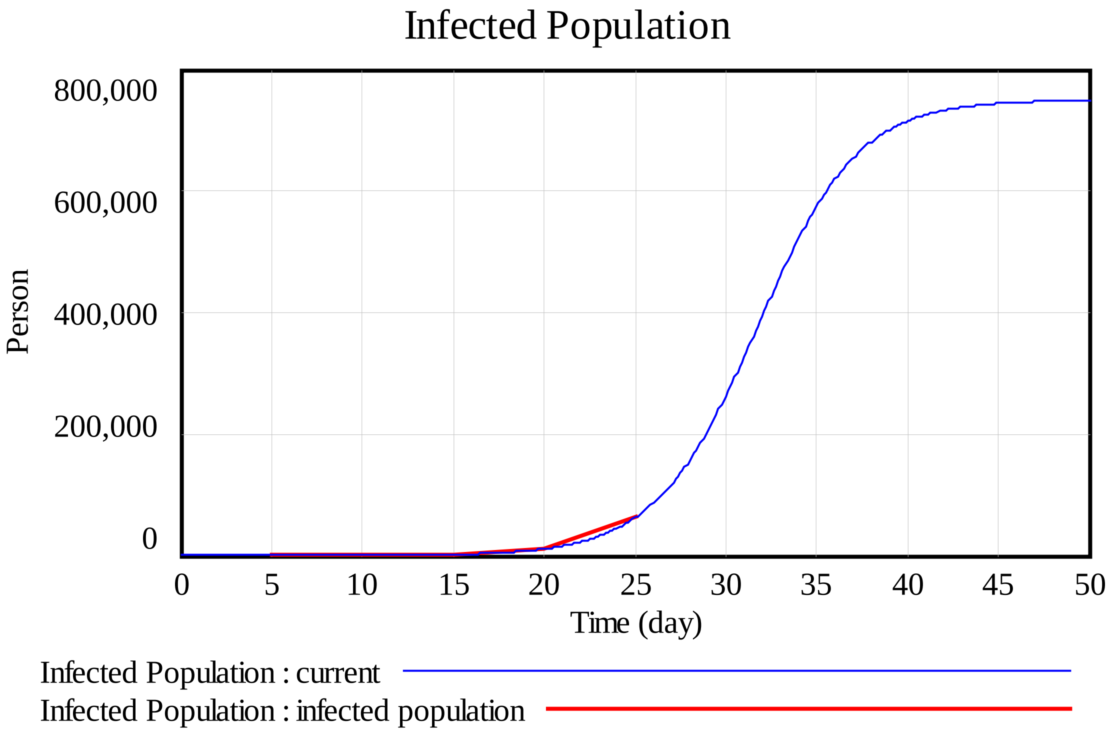

Figure 4: The cumulative model variable (infected population) together with reported data

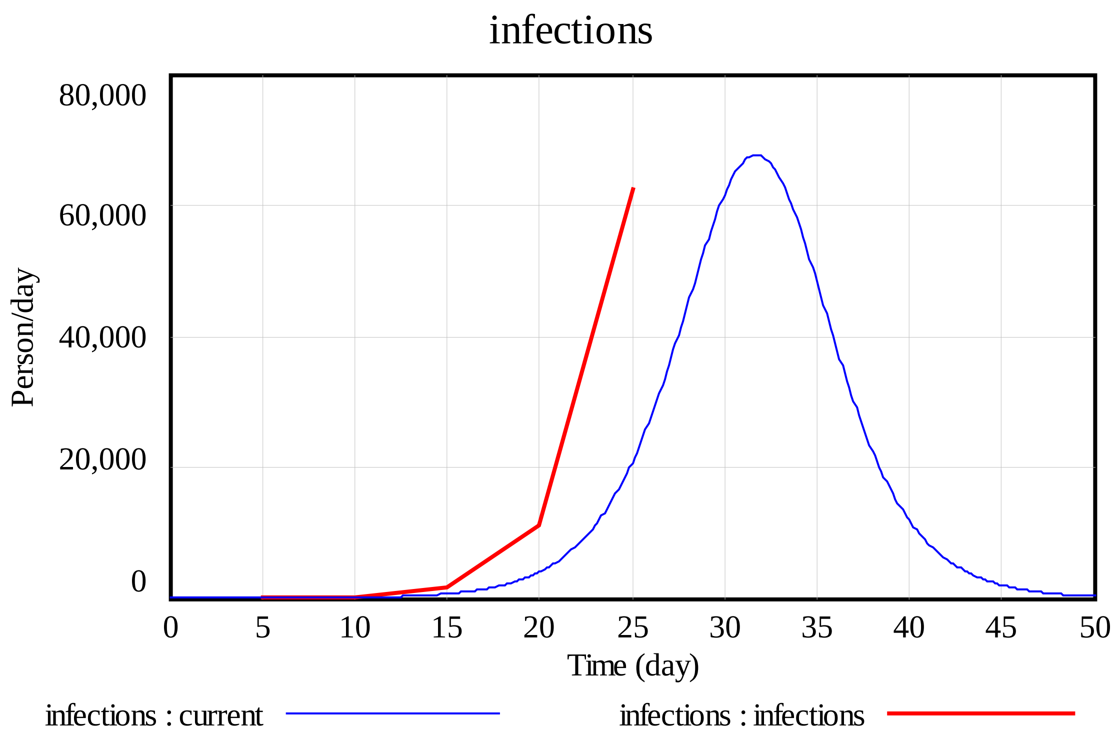

Figure 5: The corresponding fit to Infections is poor

As an aside it is worth pointing out that this model, with suitable changes to the variable names and the time constants involved, could equally represent the diffusion of a new product into a virgin market. In systems terms the structures are equivalent. The *fraction infected from contact* is the same as, say, the fraction reached by word of mouth or advertising and the *rate that people contact other people* is a measure of the potential interactions at which new products might be mentioned amongst the members of the relevant market segment. *Infected population* is equivalent to customer base, the number of adopters of the relevant product. So it is possible to shed light on important real-world marketing parameters through a calibration optimisation of models of this general structure.

3. Optimisation of performance (policy optimisation)

An example model is to be used to illustrate the process of optimisation to improve the performance of the system and this is illustrated in figure 6. It concerns the service requirements which can arise following the sale of a durable good. These items are typically sold with a 12-month warranty and during this time the vendor is obliged to offer service if a customer calls for it. In this particular case the vendor is not being responsive in terms of staffing the service section. The result is that as sales grow the increasing number of service requests is putting pressure on the service personnel.

The delay in responding to service calls also increases and the effect of this is that future sales are depressed because of the vendor's acquired reputation for poor service response. The basic behaviour mode is overshoot and collapse.

In the model depicted in figure 6, the growth process is achieved by a RAMP function which causes sales of the good to increase linearly by 20 units per month from a base of 500 units per month.

The payoff function is restricted to the variable *Sales*. However, this need not be the case. Where a number of variables might be options in a payoff function, it is possible to assign weights to each such that the sum of the weights is 1.0 (or 100). The optimisation process will then proceed with the software accumulating a weighted payoff which it will attempt to maximise. Weights are positive when more is better and negative when less is better.

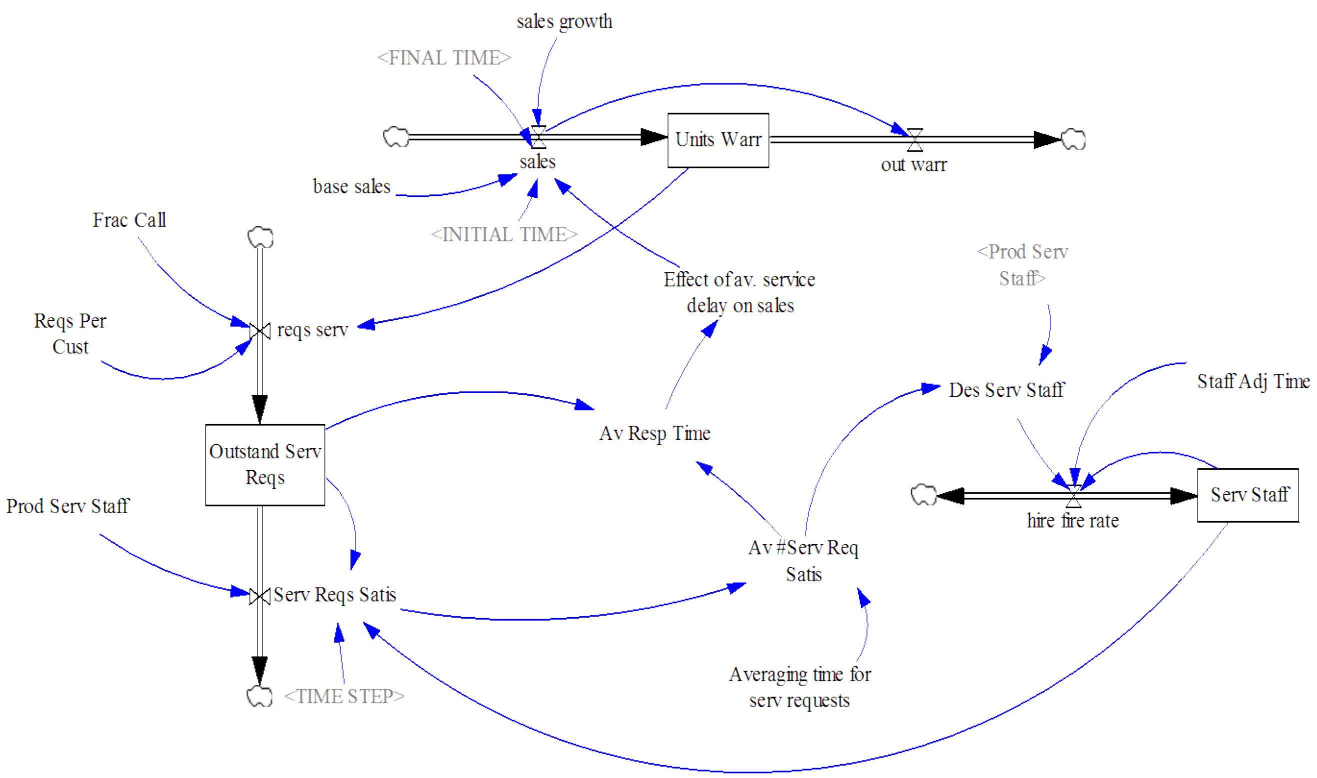

Figure 6: Model of service delays for durable goods under warranty

Policy experiment No. 1

Here it is decided to try to improve the productivity of the service staff. Currently they manage, on average, to respond to 120 calls per operative per month. It may be an option to improve their productivity by, say, providing them with hand-held devices which direct each operative from one call to the next – calls which may have arisen since setting out from their base. In this way their call routing is improved.

The optimisation parameter is *Prod Serv Staff* and we select an upper limit for the search range of 240 calls per person per month. The chosen performance variable is *Sales* since we wish to maximise this – or at least not have it overly depressed by poor response times. The results are shown in Table 5. We see that the payoff is increased and that the optimum productivity is a modest increase of 2.6 requests per month, on average. This should be easily achievable and perhaps without expenditure on high-tech devices. The graphical output for sales is shown in figure 7.

For comparison, the effect of increasing the productivity to as high as 150 calls per month, on average, is also shown. This would represent an increase of 25% and would be much more difficult to accomplish. Here the benefit of optimisation is highlighted. A modest increase in productivity returns a visibly improved sales performance (although the basic behaviour mode is unchanged), whilst a much greater productivity increase offers little extra benefit for the effort and cost involved in improving productivity.

Table 5: Optimisation results for the productivity of the service staff

Initial point of search.
Prod Serv Staff = 120.
Simulations = 1.
Pass = 0.
Payoff = 27743.5.
---------------------------------.
Maximum payoff found at:.
*Prod Serv Staff = 122.647.
Simulations = 27.
Pass = 3.
Payoff = 29915.
---------------------------------.

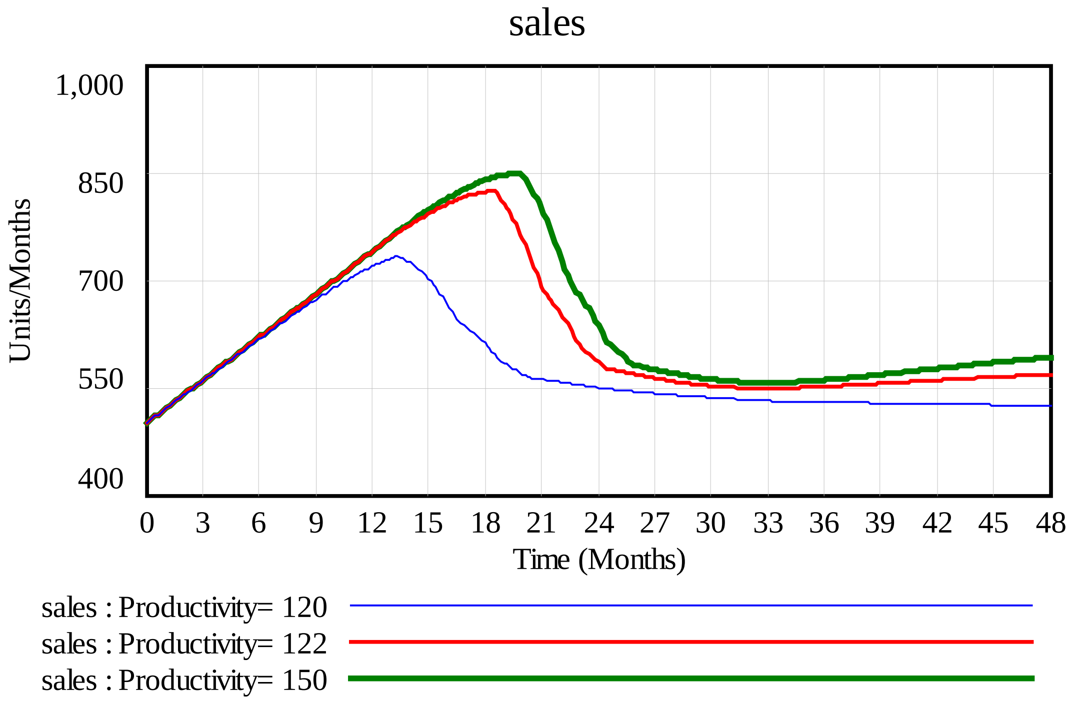

Figure 7: Plots of sales achieved for differing productivities

Policy experiment No. 2
Another approach to policy optimisation involves the use of a zero-one parameter which has the effect of either including or excluding an influence on policy. Suppose it was thought that the quantity of product units in warranty should exert an influence on the numbers of service personnel hired (or fired). The equation for the desired number of service staff (*Des Serv Staff*) can be expressed as:

Des Serv Staff= "Av #Serv Req Satis"/Prod Serv Staff* trigger+ ("Av #Serv Req Satis"/Prod Serv Staff)*(Units Warr/initial units in warranty)*(1-trigger)
(Units: Persons)

The *trigger* variable is initially set to 1.0 and so the more sophisticated policy is not active. The optimisation run results are shown in Table 6. Clearly there is benefit from including the more sophisticated policy which takes into account the current numbers of product units in warranty.

Table 6: Optimisation results from selection of policy drivers

Initial point of search.
trigger = 1.
Simulations = 1.
Pass = 0.
Payoff = 27743.5.
--------------------------------.

Maximum payoff found at:.
*trigger = 0.
Simulations = 13.
Pass = 3.
Payoff = 47090.9.
---------------------------------.

The graphical output is unequivocal (figure 8). Sales are continuously increasing when the recruitment policy for service personnel takes into account the number of product units in warranty. The depressive effect on sales of poor service performance is non-existent.

Whilst this might seem an obvious policy it is surprising how easily the naïve alternative might be accepted without question. The number of calls a typical operative can manage each month is well known along with the (historical) number of service requests satisfied. Hence, the desired number of staff is more or less fixed. This comes undone when there is a growth in the number of products sold. In this different environment such a simplistic policy can, as shown, lead to overshoot and collapse. Notice needs to be taken of the changing number of product units in warranty in order that a more effective system performance is achieved.

The above experiments are illustrative only and there is no intention of over-working a simple teaching model in order to uncover an ideal policy. In the case of policy optimisation a wide range of possible alternatives exists. Indeed, a process of learning naturally arises through carrying out repeated optimisation experiments with the model (Coyle, 1996).

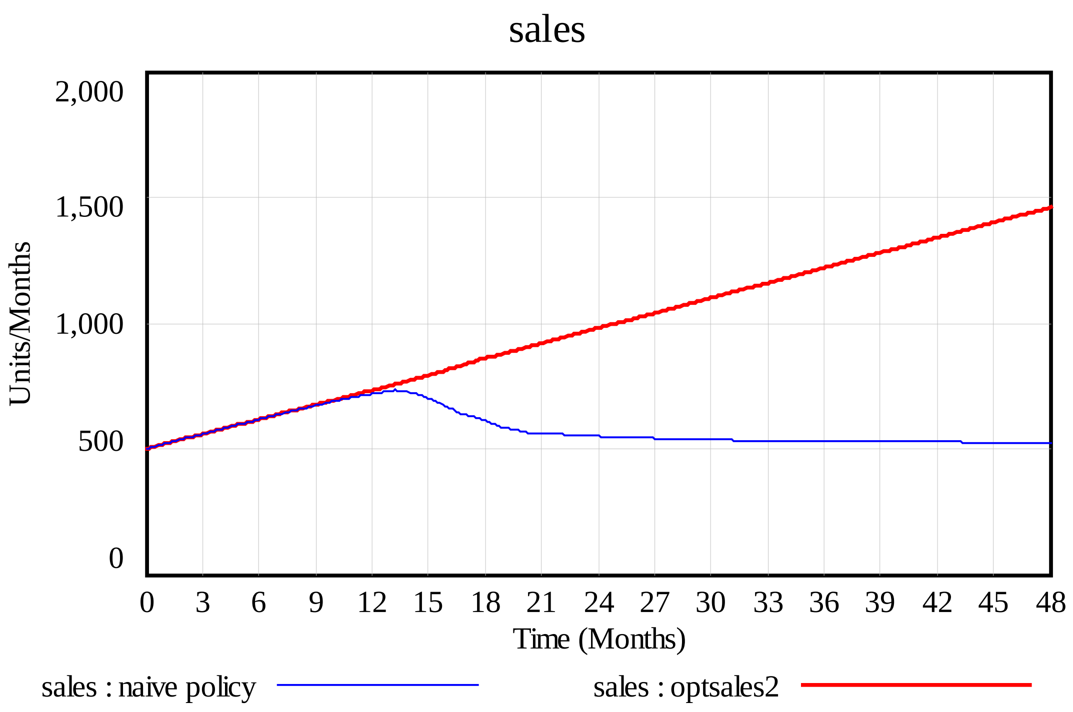

Figure 8: Comparison of sales from two different policy drivers

4. Examples of SD optimisation reported in the literature

Amongst the earliest work in this area the writings of Keloharju are worthy of mention. He contributed a number of papers on the topic in the pages of *Dynamica*. See, for example, Keloharju (1977). His work brought the concept to prominence but he did not employ the method on anything other than problems described in text books or postulated by himself. For instance, an application of optimisation to the project model contained in Richardson and Pugh's (1981) text is contained in Keloharju and Wolstenholme (1989). A statement of the method together with some textbook examples is also available (Keloharju and Wolstenholme, 1988). Additionally, an overview of the methods and their deployment on textbook examples has been contributed by the current author (Dangerfield and Roberts, 1996a). Finally, there is an example of optimisation applied to defence analysis. Again though it is a standard defence model – the armoured advance model – rather than any real-world study (Wolstenholme and Al-Alusi, 1987)

Retaining the emphasis on textbook problems for the moment, Duggan (2005) employs Coyle's (1996) model of the *Domestic Manufacturing Company* to illustrate the methods of multi-objective optimisation – an advance over standard SD optimisation with its single objective function. The concept of multiple objectives arises from multi-criteria decision-making where a situation can be judged on more than one performance metric. While a multi-objective payoff function can be formulated using a set of weights, it is argued that the selection of the weights is very individual-specific. The multi-objective approach – underpinned by the methods of genetic algorithms – rests upon determining a Pareto-optimal situation, defined as one where no improvement is possible without making some other aspect worse. In other words the method strives for an optimal solution which is not dominated by any other solution. The author demonstrates the approach combining two objectives in the model: one for the differences between desired stock and actual and another between desired backlog and actual.

In terms of applications to real-world problems, the current author has also used the methods of optimisation in research conducted in connection with modelling the epidemiology of HIV/AIDS. Fitting a model of AIDS spread to data was carried out for a number of European countries (Dangerfield and Roberts, 1994; Dangerfield and Roberts, 1996b). The optimised parameters furnished support for some of the features of AIDS epidemiology which, at the time, were being uncovered by other branches of science. For example, the optimised output revealed that a U-shaped profile of infectiousness in a host was necessary in order to achieve a best fit to data on new AIDS cases. This infectiousness profile was also evidenced by virologists who had analysed patients' blood and other secretions on a longitudinal basis.

Within this strand of research a much more complex optimisation was performed using American data on transfusion-associated AIDS cases (Dangerfield and Roberts,1999). The purpose here was to estimate the parameters for a number of plausible statistical HIV incubation distributions. Given the nature of the data the point of infection could be quite accurately determined, but two difficulties were evident: the data were right-censored and the number receiving infected transfusions in each quarter was unknown. However, the SD optimisation could estimate this number as part of the process, in addition to estimating parameters of the incubation distribution. The best fit was found to be a three-stage distribution similar to the gamma and one which accorded with the high-low-high U-shaped infectiousness profile which was receiving support from a number of sources.

In the marketing domain Graham and Ariza (2003) carried out an optimisation on a system dynamics model which was designed to shed light on the allocations to make from a marketing budget in a high-tech client firm. Assuming the budget was fixed, the task was to optimise the allocations across more than 90 'buckets' – combinations of product lines, marketing channels and types of marketing. However, these were not discrete: advertising on one product line might have crossover effects on another and the impacts could propagate over a period of time. One major conclusion for this firm was that the advertising allocation should be increased markedly. In general intuitive allocations were shown to fall short of the ideal: they were directionally correct but magnitudes fell short often by factors of three or four.

5. Future directions in SD optimisation

A primary aim must be to see more published work which describes optimisation studies carried out on real-world SD applications. There may be frequent use of optimisation in consulting assignments but such activities are rarely published. The references herein suggest that, thus far, outside of unpublished work, the number may be three at most. Whilst software requirements may have inhibited use of SD optimisation in the past, there are now no computational barriers to its use and it is to be hoped that in future this quite powerful analytical tool in SD will feature in more application studies.

An advance in the methodology itself has been developed by Duggan (2008) and this is a promising pointer for the future. Based on genetic algorithms, it is best suited to the class of SD problems that are agent-based and this highlights a slight limitation. Traditional optimisation takes the policy equations as given and explores the parameter space to determine an optimal policy. Instead he has offered an approach which searches over both parameter space and policy (strategy). Theoretically there is no limit to the number of strategies which can be evaluated in this approach, although the user has to define a set in advance of the runs. Under a conventional optimisation approach a limited tilt at this is possible using the zero-one parameter method suggested above, although this would restrict the enumerated strategies to two only. Duggan demonstrates the new approach using a classic SD problem: the four agent beer-game. We await its use in a real-world application.

6. References

Coyle RG (1996) *System Dynamics Modelling: a practical approach*, Chapman & Hall: London.

Dangerfield BC and Roberts CA (1994) Fitting a model of the Spread of AIDS to data from five European countries. In: *O.R. Work in HIV/AIDS* (2nd Edition), Operational Research Society, Birmingham, 7-13.

Dangerfield BC and Roberts CA (1996a) An Overview of Strategy and Tactics in System Dynamics Optimisation, *Journal of the Operational Research Society*, 47:3 405-423.

Dangerfield BC and Roberts CA (1996b) Relating a transmission model of AIDS spread to data: some international comparisons. In: V. Isham and G. Medley Eds., *Models for Infectious Human Diseases: Their Structure and Relation to Data*, Cambridge University Press, 473-476.

Dangerfield BC and Roberts CA (1999) Optimisation as a statistical estimation tool: an example in estimating the AIDS treatment-free incubation period distribution, *System Dynamics Review*, 15:3, 273-291.

Duggan J (2005) Using Multiple Objective Optimisation to generate policy insights for system dynamics models, *Proceedings of the International System Dynamics Conference,* Boston. System Dynamics Society. (CD-ROM).

Duggan J (2008) Equation-based policy optimisation for agent-oriented system dynamics models, *System Dynamics Review*, 24:1, 97-118.

Graham AK and Ariza CA (2003) Dynamic, hard and strategic questions: using optimisation to answer a marketing resource allocation question, *System Dynamics Review*, 19:1, 27-46.

Keloharju R (1977) Multi-objective decision models in System Dynamics, *Dynamica*, 3:1 3-13 and 3:2, 45-55.

Keloharju R and Wolstenholme EF (1988) The Basic Concepts of System Dynamics Optimisation. *Systems Practice,* 1: 65-86.

Keloharju R and Wolstenholme EF (1989) A Case Study in System Dynamics Optimisation, *Journal of the Operational Research Society,* 40:3, 221-230.

Meadows DM and Robinson JM (1985) *The Electronic Oracle*, Wiley: Chichester. (Now available from the System Dynamics Society, Albany NY.)

Richardson GP and Pugh AL (1981) *An Introduction to System Dynamics Modelling with* DYNAMO, MIT Press, MA. (Now available from Pegasus Communications, Waltham, MA.)

Wolstenholme EF and Al-Alusi AS (1987) System dynamics and heuristic optimisation in defence analysis, *System Dynamics Review*, 3:2 102-115.