A patient-level simulation model to estimate the costs and outcomes for various methods of screening for *Chlamydia trachomatis*

Pelham Barton, Suzanne Robinson, Tracy Roberts
Health Economics Facility, University of Birmingham

Address for correspondence:
Pelham Barton
HSMC
Park House
40 Edgbaston Park Road
Birmingham
B15 2RT

Abstract

Chlamydia trachomatis is the most common sexually transmitted infection in the UK. While usually asymptomatic, it can lead to serious complications. For women complications include pelvic inflammatory disease, infertility, and ectopic pregnancy.

Modelling the effects of screening for Chlamydia requires a modelling technique which respects the natural history of the infection. Models which do not allow for interaction between individuals omit two important effects. Individuals who are successfully treated for Chlamydia will not thereafter pass the infection on to others. On the other hand, if current sexual partners of infected individuals are not also treated, the treatment effect is likely to be lost as treated individuals are re-infected.

The issue of partner notification is thus seen to be an important part of any programme to reduce the prevalence of Chlamydia. Our model builds on the previous model by Mirjam Kretzschmar. It follows a hypothetical population through formation and dissolution of partnerships, transmission of Chlamydia, testing and treatment, partner notification, and long-term consequences.

Acknowledgements
This work forms part of the HTA project reference number 97/32/31 "A study to evaluate the most cost effective way to screen for Chlamydia trachomatis genital tract infection and reduce its prevalence and associated burden of disease". The study is known as the Chlamydia Screening Study (ClaSS).

We would like to thank other members of the ClaSS project team, and especially our external advisors Mirjam Kretzchmar, Liam Toohill and Robert Welte. We accept responsibility for all statements made in this paper.

Introduction

Chlamydia trachomatis is the most common sexually transmitted infection in the UK. While usually asymptomatic, it can lead to serious complications. For women complications include pelvic inflammatory disease, infertility, and ectopic pregnancy. This paper reports the development of a model designed to estimate the costs and consequences associated with various methods of screening for the *Chlamydia trachomatis* infection. The modelling work forms part of the Chlamydia Screening Study (ClaSS).

Modelling the effects of screening for Chlamydia requires a technique which respects the natural history of the infection. Models which do not allow for interaction between individuals omit two important effects. Individuals who are successfully treated for Chlamydia will not thereafter pass the infection on to others. On the other hand, if current sexual partners of infected individuals are not also treated, the treatment effect is likely to be lost as treated individuals are re-infected. Since the two biases in "static models" are in opposite directions, the result from any such model cannot be used as a basis for any rational decision making. Therefore, the only question in selecting a modelling approach is whether to use an aggregation approach akin to System Dynamics (Townshend and Turner, 2000), or an individual-based approach akin to Discrete Event Simulation (Kretzschmar *et al*, 1996).

The ClaSS study collected individual level data on both screening and partner notification which led us to the approach that would best utilise these data, namely Discrete Event Simulation. For example, the ClaSS project is comparing various forms of partner notification. In the System Dynamics approach, this could only be done by estimating the average effects, whereas in the Discrete Event Simulation approach, records can be kept of partners attached to specific individuals.

The basic model structure

The ClaSS model is a new model based on the framework by Kretzschmar and colleagues (1996). A population is simulated over time, with individual characteristics changing as necessary on a daily basis. Although in the implementation these are mixed together, the features of the model may conveniently be divided into the following sections:

- ageing and replacement;
- partnership formation and dissolution;
- Chlamydia transmission and progression;
- testing and treatment;
- partner notification.

The description which follows is of the current state of the model. In particular, this applies to any reference to a fixed probability: the model structure could easily

accommodate varying such probabilities by any personal characteristic(s) in any way that can be specified.

Ageing and replacement

The initial population consists of a number of virtual individuals with ages drawn from a uniform distribution between lower and upper limits. By default, the ages range from 15 to 65, and individuals are equally likely to be male or female. As the model runs, individuals in the model die (from "other causes") in line with standard UK life tables, and new individuals at age 15 are added to the model. Chlamydia related death is considered negligible and therefore not included.

Partnership formation and dissolution

As with previous models (Kretzschmar *et al*, 1996; Townshend and Turner, 2000), only heterosexual partnerships are considered. The initial population does not contain any partnerships. During the running of the model, new partnerships form and are dissolved. Individuals are divided into two risk groups, a "core group" of highly promiscuous individuals, and the rest. Individuals' propensity to form new partnerships is determined in part by their risk group and existing partner status, also on the age difference between the prospective partners. A further characteristic of the partnership is the frequency of sexual contact.

It is assumed that the probability of any partnership ending (for reasons other than death) on a given day is independent of other partnerships in which either partner may be involved.

Chlamydia transmission and progression

An individual's Chlamydia status is represented by a number from the following list:

0: no Chlamydia;
1: latent Chlamydia;
2: asymptomatic Chlamydia;
3: symptomatic Chlamydia.

(As the model develops, further numbers will be used to represent the various sequelae of Chlamydia infection.)

Initially, a small proportion of the population is infected with Chlamydia. If one member of a partnership is infected, but the other is not, then there is a risk of transmission on each sexual contact. This is combined with the frequency of sexual contact in the partnership to give the probability of transmission on any day. Different figures are used for male-to-female and female-to-male transmission.

An individual who is infected enters a period with latent Chlamydia, which is assumed to last a fixed number of days (different for male and female), and then may

become asymptomatic or symptomatic. Individuals with latent Chlamydia cannot pass on the infection.

For any infected individual, there is a sex-dependent fixed probability of spontaneous recovery on any given day. Symptomatic individuals also have a fixed probability of seeking treatment.

Testing and treatment

As well as a "no screening option", three different screening policies are currently programmed.

Method 1 – Single Pulse. Screening is assumed to start on a given date after the start of running the model. Once the screening programme has started, females within a certain age range (by default 16 to 24 inclusive) are eligible for selection as index cases. Women who have previously been involved with the screening programme are not eligible. On each day, each eligible woman has a fixed probability of being selected.

Method 2 – Screen by Ages (female). Women are selected for screening on reaching any of a given set of ages. (By default, these run from 16 to 27 inclusive.)

Method 3 – Screen by Ages (all). Men and women are selected for screening on reaching any of a given set of ages.

Those selected may then comply with the screening test. Those who comply receive a screening result at a fixed delay after testing. Allowance is made for sensitivity and specificity of the test used to be below 100%. Different forms of specimen collection can be used for males and females: accordingly, the sensitivity and specificity are sex-dependent. Those who screen negative are not treated, but may be called for screening in the future. Those who test positive will be treated.

Partner notification

Anyone treated for Chlamydia, either a symptomatic individual requesting treatment or one found through screening, will be asked to notify partners. Those who comply with partner notification will inform all current partners and former partners within a specified time interval,(for simplicity, it is assumed that notification here is either zero or complete). Allowance is made in the model for different rates of compliance between symptomatic cases and those found through screening. Each partner individually may then comply with a request to attend for treatment. Any partner attending will be treated without testing.

Personal information maintained in the model

The assumption that "other causes" death is completely independent of partnership and Chlamydia status allows each individual's date of death to be sampled when that individual is introduced into the model. Similarly, the date at which a partnership is due to end can be sampled when the partnership is formed.

Each individual in the model carries a person record, containing the following items:

- whether individual is still alive,
- sex of individual,
- age, measured in days,
- risk group,
- number of current partners,
- date of death,
- Chlamydia status,
- date of last change of Chlamydia status,
- Chlamydia status at time of screening test,
- date at which screening result is due,
- date of last encounter with the screening and treating programme,
- result of most recent Chlamydia test,
- date on which to respond to partner notification,
- list of current partners, with date at which partnership started, date at which partnership will end, and frequency of sexual contact during the partnership,
- list of former partners, with details as for current partners.

The list of current partners is stored as a fixed-length array of integers, with zeroes filling the gaps. Dates of starting and ending and sexual contact rate are stored in parallel arrays. The list of former partners is stored as a string, and grows in length with time. Whenever a partnership is terminated, a 24-character string is created, giving identity number of former partner, a code letter for reason for termination (death or otherwise), and the start and end dates of the partnership. This is appended to the list of former partners.

Implementation of the model

The ClaSS model is coded in Borland Delphi, which is based on Pascal. It is compiled into an executable file, which can be run under any suitably recent version of Windows, without the need for any special software or licence. The program takes input from various data files, which must be in the same directory as the program file.

Running the model

When the program is loaded, the screen as shown in Figure 1 appears.

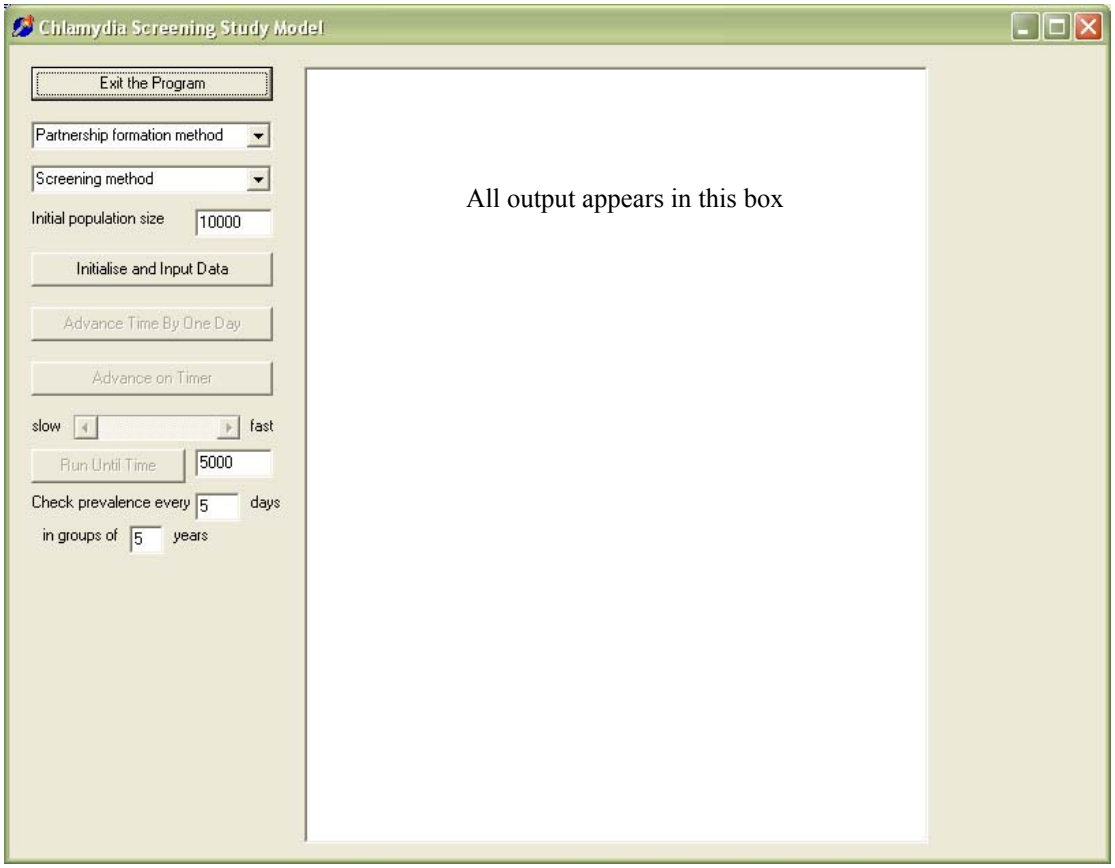

Figure 1. Initial screen for the ClaSS model

To start running the model, it is necessary to click on the button "Initialise and Input Data". Before doing this, the initial population size may be varied between a minimum of 2000 and a maximum of 100,000. (Values below 2000 are changed to 2000; values above 100,000 to 100,000.) Note that the maximum total population currently allowed is 100,000: if the initial population is set at 100,000, no new individuals will be added to the model. It is also possible to select from a choice of different methods of modelling partnership formation. At present, four are available:

- female first then male (the default option);
- male first then female;
- either first;
- symmetric.

The default method works by considering in turn each living female. According to her age, risk group, and current partnership status, she has a certain probability of forming a new partnership on any day. If she is selected to form a new partnership, then living males are successively considered at random until one is accepted. The probability of acceptance depends on age difference, and the propensity of the male to form a new partnership. Once a male has been accepted, the length of the partnership is determined. Each partner's list of regular partners is then updated accordingly.

The second method works in the same way, except that the roles of males and females are reversed. The third method is a combination of the first two.

The fourth method ("symmetric") is slightly different. Firstly the number of new partnerships to be formed on any day is determined. Then new partnerships are formed as follows. A female and a male are selected at random. The probability of them forming a new partnership depends on each one's propensity to form a new partnership at that time, and on compatibility rules based on age. If the new partnership is accepted, then its length is determined; otherwise, both potential partners are discarded and a new pair drawn. This process is repeated until the required number of new partnerships has been formed. Note that this method requires many more draws from the random number generator than the other three.

Other methods may be implemented in later versions of the model.

The options for screening are:

- No screening (the default option);
- Single pulse;
- Select by age (females);
- Select by age (all).

When the user clicks on the "Initialise and Input Data" button, the data files are read, and an initial population is generated. When this has been completed, the screen changes to that shown in Figure 2.

The two buttons which run the model are now enabled, as are a variety of display options. When the model is running, it operates in cycles of one day. The button "Advance Time By One Day" causes a single cycle to be run, and the display to be updated, after which the program waits for further action by the user. The button "Advance on Timer" starts continuous running, which continues until the same button (now labelled "Halt Timer") is pressed. The button "Run Until Time" allows the model to be run at maximum speed up to a specific time.

What actually happens when the program is running under "Advance on Timer" depends on the position of the slider immediately below that button. When it is at the end marked "slow", the display is updated and there is a wait of approximately two seconds (actual time) after each simulated day's running. As the slider is moved to the right, the wait is halved until it reaches a negligible amount. Moving further to the right, the number of simulated days running between each screen update is doubled up to a maximum of 16. In the fastest running mode, the screen updating is sufficient to show that the program is still working, but makes very little difference to the total running time.

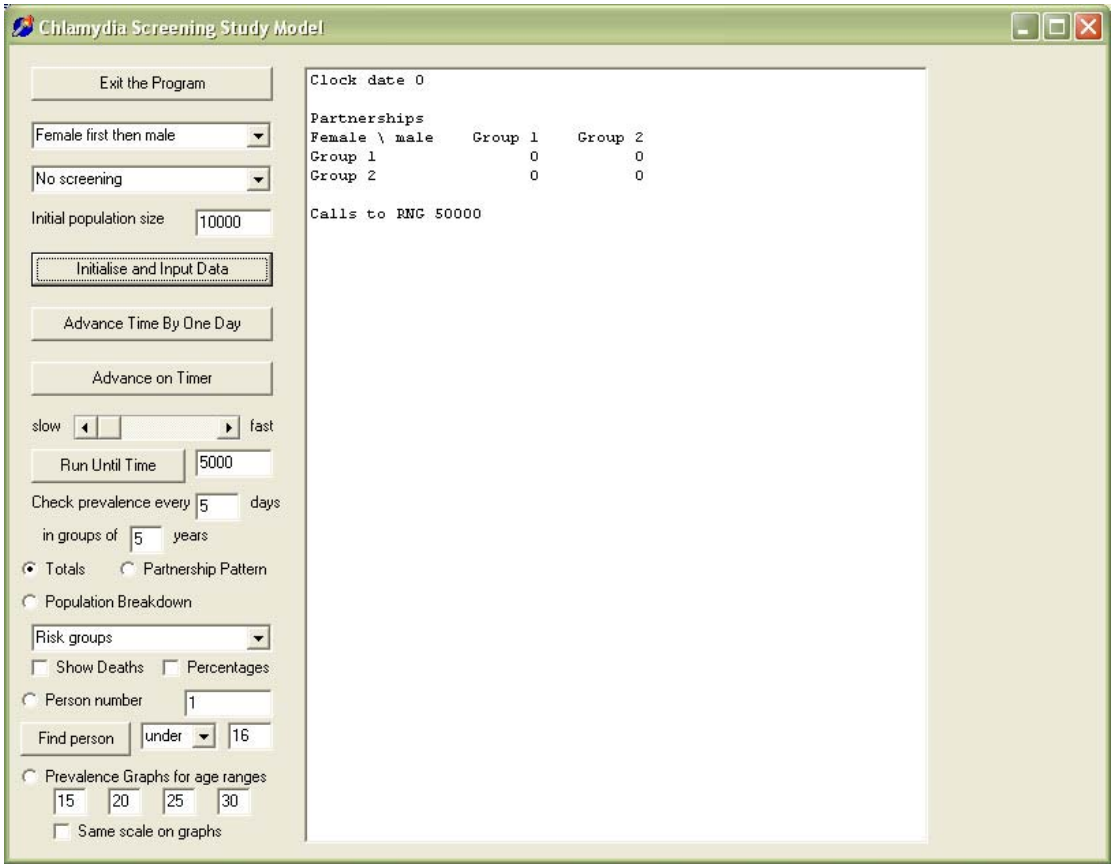

Figure 2. The ClaSS model ready for running

The display options can be selected by use of the "radio buttons". Each option gives the clock date (in days from the start). The remainder of the display depends on the option selected as follows:

Totals

This gives the total number of partnerships, classified according to the risk group of the partners. This output will be used to calibrate the model to ensure the appropriate mixing within risk groups. It also shows the total number of calls to the random number generator. The random number generator currently used has a cycle length of approximately $2^{62} \approx 4.6 \times 10^{18}$.

Partnership Pattern

This shows the current pattern of partnerships within the model. A typical screen is shown in Figure 3. The display area is made up of a grid of squares for each possible combination of ages of the female and male partners. The more partnerships there are for any age combination, the darker the colour of the corresponding square. For example, in the diagram, there are (at this particular time for this particular run of the model) partnerships involving 15-year-old females with males of ages 17, 18, 19, 20 21 and 26, but no other ages. The shading is scaled so that black represents the largest number of partnerships currently in existence for any age combination: in this case,

there are at most seven partnerships for any such combination, and there are seven partnerships between 17-year-old females and 16-year-old males.

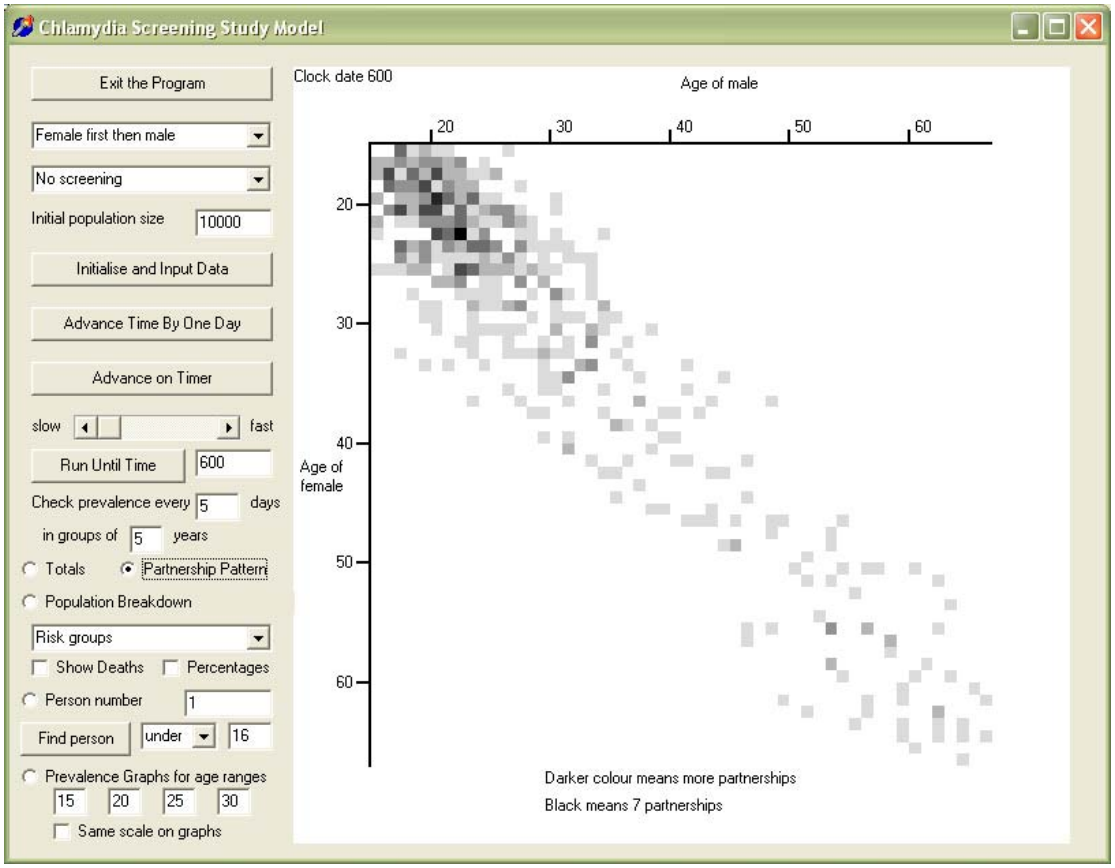

Figure 3. Example of partnership patterns.

Population Breakdown

This shows a breakdown of the current population within the model. The population is divided into single year age groups. The total number (or percentage) of females and males in each age group is shown, classified by any one of: risk group, partner status or Chlamydia status. This output only makes sense in colour, and has been omitted from this paper.

Person Number

This shows the individual details for the person numbered in the edit box immediately to the right of the radio button. The number must be within the range from 1 up to the total number of individuals within the model. Numbers outside that range are replaced by numbers within the range. It is also possible to search through the population for a person above or below a given age. The example shown in Figure 4 is of a 22-year-old male in risk group 1, who has one current partner and has had 11 previous partners within the running of the model. Note that the date of his death, and the date of ending of the current partnership are shown. The identity numbers of his current and former partners are shown: among other things, this display allows checking that

partnerships are consistent. For example, the person record for number 1143 must show the same current partnership as is shown here.

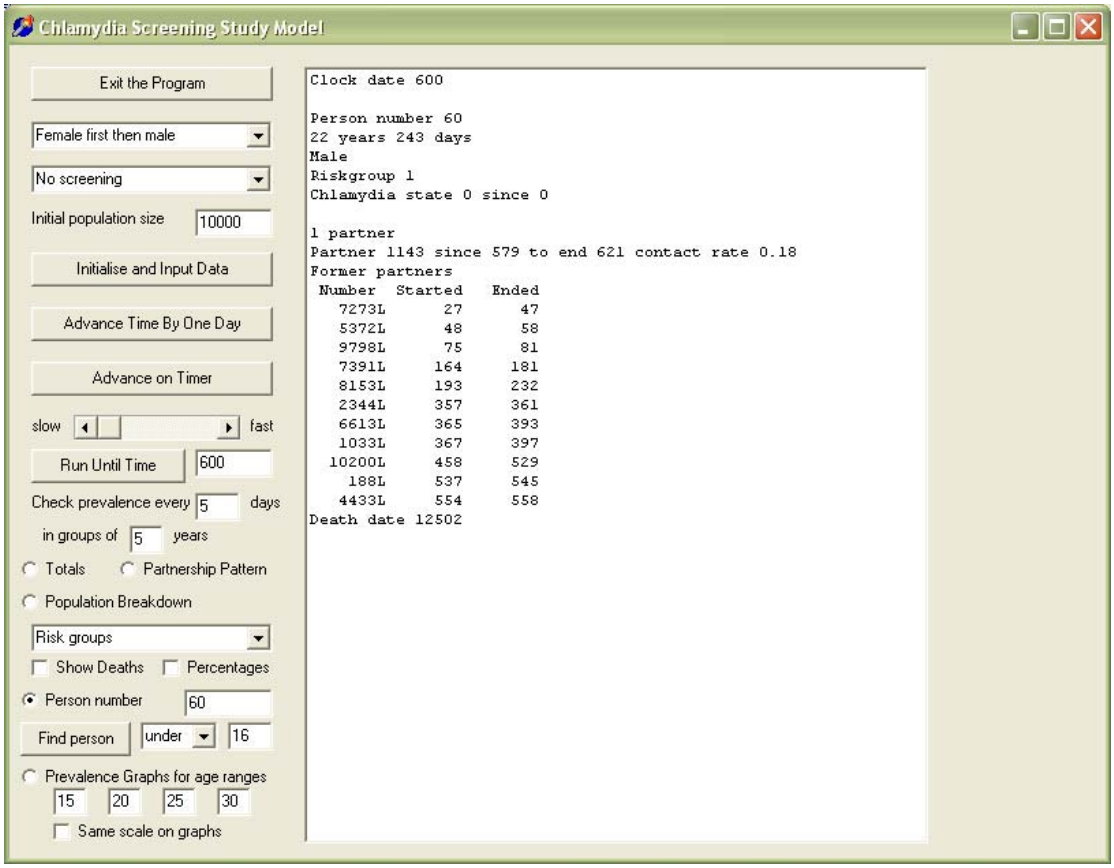

Figure 4. Individual person record.

Prevalence Graphs

The examples shown in Figure 5 are in 5-year age groups, with the data collected every five days. (These are the default parameters – any changes must be made before running the model.) The horizontal scale shows time in days, while the vertical scale shows the prevalence. By default, the top of the vertical scale is the lowest multiple of 0.01. It is possible to set the four graphs to have equal scales (to allow a direct visual comparison between the age-groups) and to set a fixed value for the top of the scale (to allow a comparison between the effects of different strategies): examples of this are shown later. Note that the lowest of the four graphs shown here (age-group 30 meaning 30 to 34) shows a stepped appearance: the top of the scale here is typically around 15 people.

For clarity with a black and white printer, the graph for females is thicker than the graph for males.

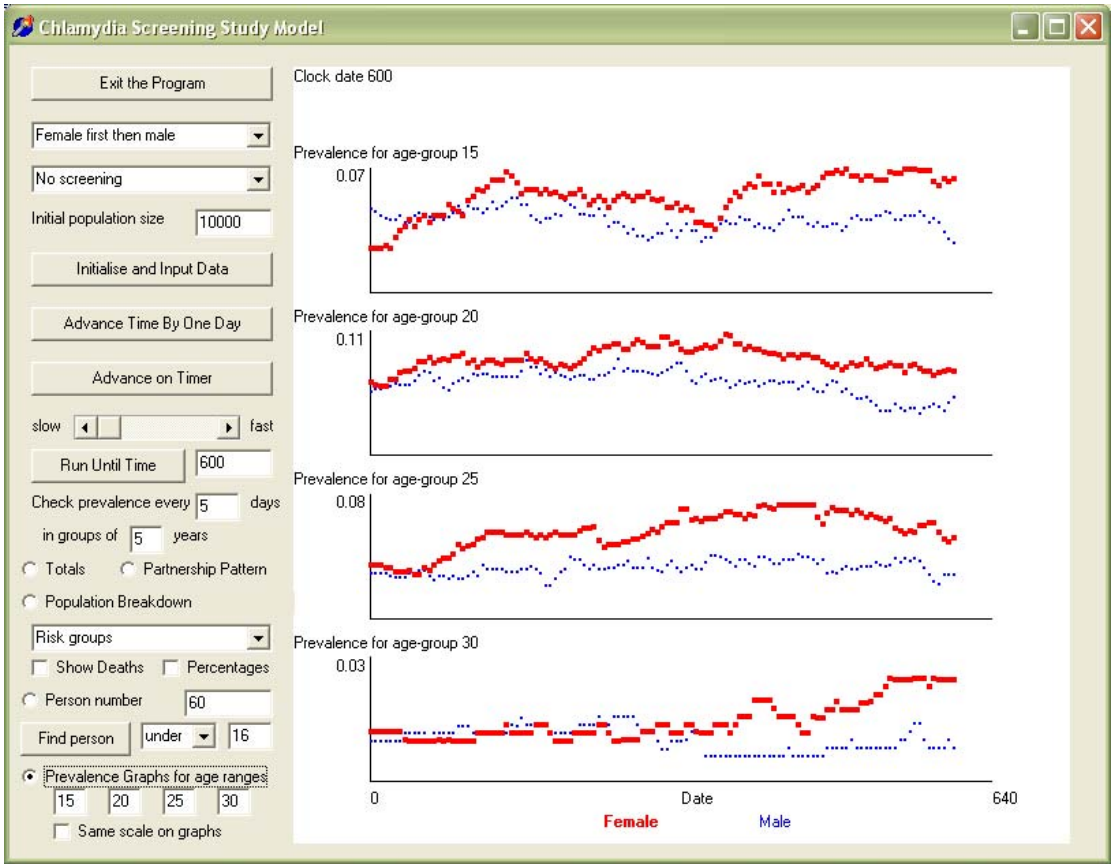

Figure 5. Prevalence graphs

Illustrative results

The data inputs for these results are still largely arbitrary: accordingly, the actual values used are not described. To obtain reasonably smooth graphs, a starting population of 50000 was used. Figure 6 shows the result of running for 10000 days with the "No screening" option. It can be seen that there is a "warm up" period, after which the prevalence continues to vary with time, but with no trend.

The steady-state prevalence with no screening should be a function of the partnership and transmission data, and independent of the starting prevalence. To test this, the model was re-run with only half the initial prevalence. The results are shown in Figure 7. As can be seen, the warm-up period appears to be slightly longer, but the final prevalences are much the same as in Figure 6.

The effect of introducing screening after the warm-up period is shown in Figure 8. There is a considerable drop in prevalence from the end of the warm-up period. The advantage of using a warm-up period is that the model can give an idea of the short-term effects of introducing a screening programme, as well as the steady state that is eventually reached under such a programme.

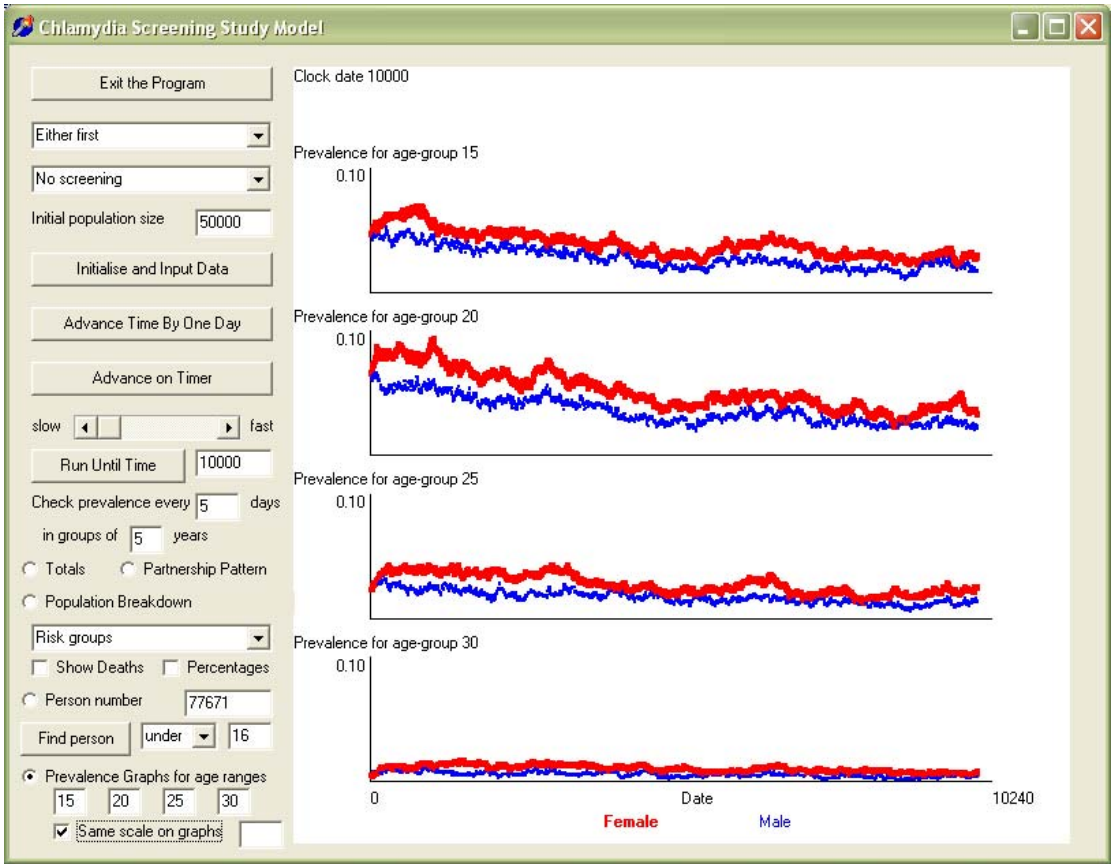

Figure 6. Illustrative result with No Screening option selected.

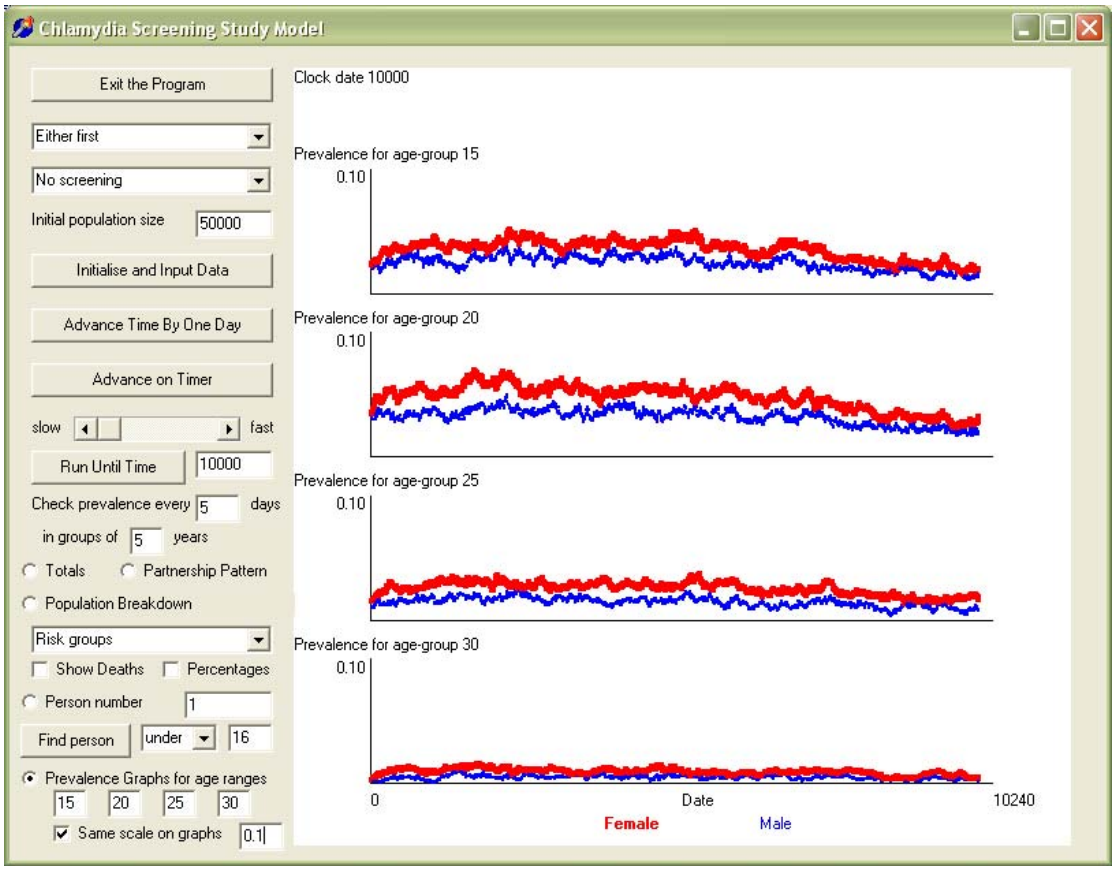

Figure 7. Starting prevalence halved.

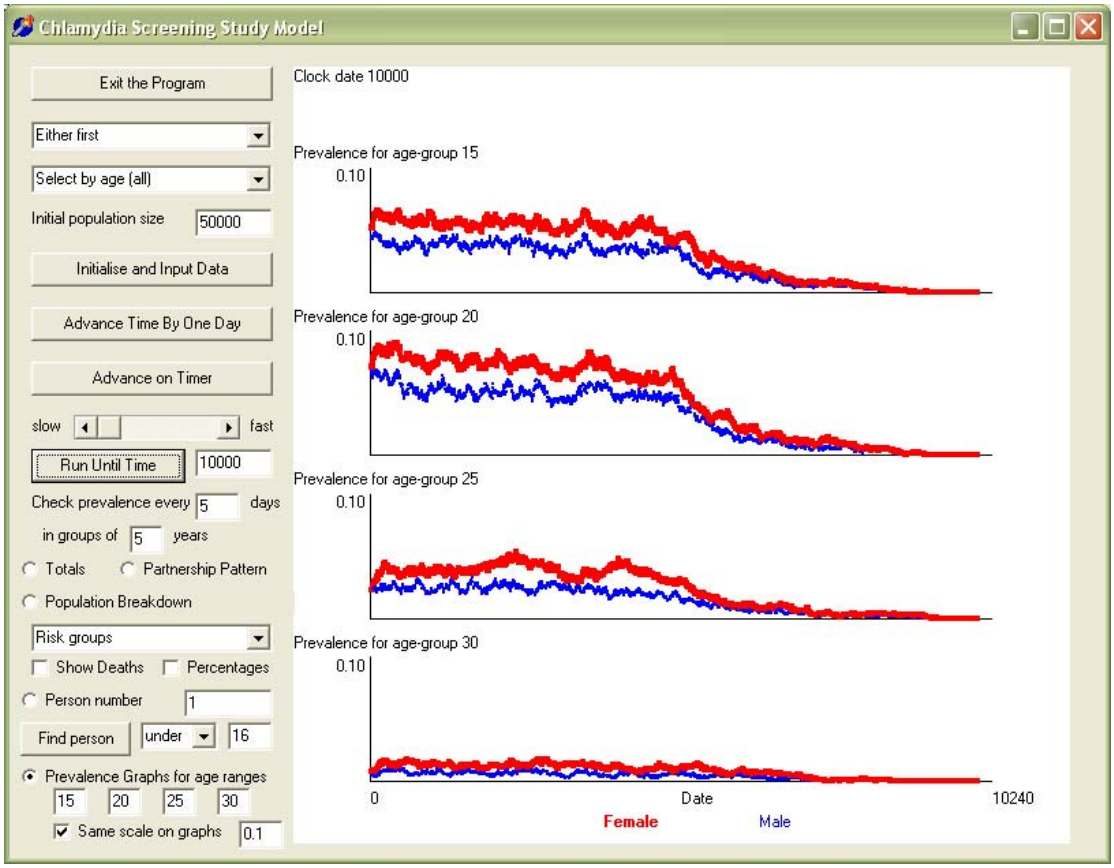

Figure 8. Screening introduced after warm-up period.

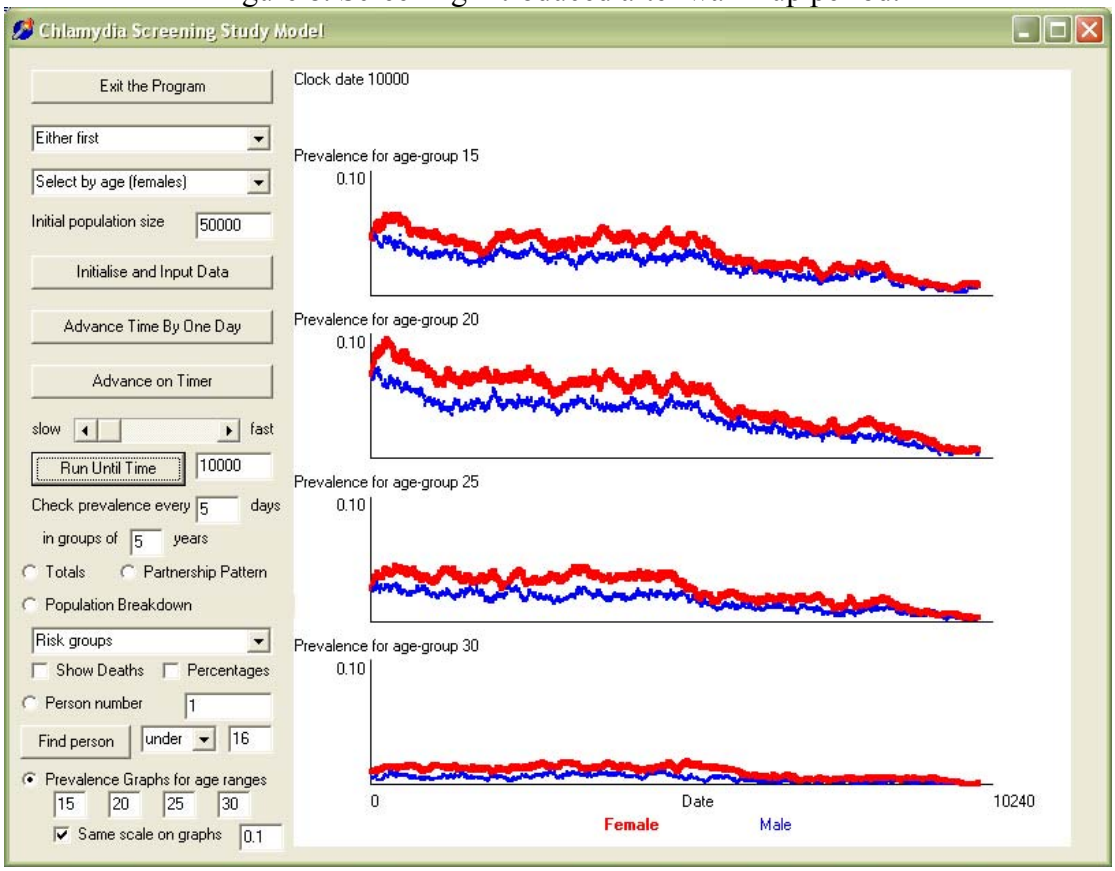

Figure 9. Screening of females only.

It has been suggested that it is sufficient to screen females only, relying on partner notification to identify infected males. Figure 9 shows the effect of such a policy. By comparison with Figure 8, screening females only appears to leave a higher prevalence of Chlamydia in the population.

Costing

Resource use (testing and treatment, etc) can be counted in the model and the total cost of running any screening programme over a period of time can be estimated. This can then feed into a cost-effectiveness analysis for the programme as a whole. Note that such calculations, comparing different strategies, include benefits to individuals who have never received treatment themselves, but whose infection has been avoided by treatment of others.

Discussion

In this paper we have presented a working model using Discrete Event Simulation that will ultimately be used to evaluate the cost effectiveness of screening for Chlamydia trachomatis. With the exception of Ktretschmar et al (2000), and Townshend & Turner (2000) previous models used for economic evaluation of Chlamydia screening have tended to take a static approach which has not respected the natural history of the infection and the interaction between individuals (Roberts et al 2003).

The model as is stands is heavily dependent on the input parameters relating to partnership formation and dissolution. While some data are available, the best way of ensuring that the inputs are reasonable is by calibration. The model inputs are adjusted until the steady state prevalence of Chlamydia shows the required age distribution. Once this has been done, the effect of imposing a screening pattern can be assessed. Since there may be many different sets of input parameters which lead to the same steady state prevalence, extensive sensitivity analysis will be required.

The model as it stands does not include sequelae of Chlamydia infection. These sequelae are to be added in a later version of the model.

References

Kretzschmar, Mirjam; Van Duynhoven, Yvonne THP; Severijnen, Anton J (1996) Modelling Prevention Strategies for Gonorrhea and Chlamydia Using Stochastic Network Simulations. *American Journal of Epidemiology* 144(3); 306-317.

Roberts, Tracy; Robinson, Suzanne; Barton, Pelham; Bryan, Stirling (2003) The *Dangers of Introspection: A Systematic Review of Economic Studies of Chlamydia Trachomatis Screening*; Presented at HESG meeting Canterbury

Townshend JRP; Turner HS (2000) Analysing the effectiveness of *Chlamydia* screening. *Journal of the Operational Research Society* 51; 812-824.