

6 Application to Uganda: Results

In this chapter we apply the model to Uganda, drawing on the empirical preparatory work of the preceding chapter. Our initial aim is simply to obtain an approximate replica of the Uganda epidemic in the form of a model fitting.

Within the constraints of its design the model allows great flexibility in allowing the various inputs to vary with time, duration and sex. In the initial application we simplify by making a number of simplifying assumptions that may be relaxed in future applications. The age schedules of fertility and male and female mortality from non-HIV-related causes, the age schedules of male and female entry into and exit from sexual activity, and the age- and duration-specific schedule for mortality from HIV-related causes will all be assumed constant over time at the values indicated in the preceding chapter.

The epidemiological parameters

We focus here on the parameters that define the epidemiology of the model: the exogenous stimuli parameter $\mathbf{i}_S(t)$ that initiates and sustains the epidemic, the time-invariant and time varying endogenous parameters \mathbf{k}_S and $\mathbf{q}(t)$. Recall that the level of the force of infection in any year, given by equation (13) of the Model Design chapter, has the form

$$\mathbf{i}_S(t) + [\mathbf{k}_S + \mathbf{q}_S(t)]p_{S'}(t)$$

where S denotes sex and S' denotes the opposite sex. Thus for males the equation is

$$\mathbf{i}_m(t) + [\mathbf{k}_m + \mathbf{q}_m(t)]p_f(t),$$

and for females it is

$$\mathbf{i}_f(t) + [\mathbf{k}_f + \mathbf{q}_f(t)]p_m(t).$$

We are of course ignoring the age patterns here, since these have been factor out into the age pattern functions $\mathbf{q}_S(t)$.

Sex differences

For initial fitting of the model we simply this set of parameters in two ways, by equating male and female levels and by imposing simply time patterns on $\mathbf{i}_S(t)$ and $\mathbf{q}(t)$. There are several obvious choices for the time pattern of $\mathbf{i}_S(t)$.

The exogamous risk parameter

For an epidemic to begin the exogenous parameter must have a positive value for some year, which we may take without loss of generality to be the first year of the projection. Assuming

$\mathbf{i}_s(1) \equiv \mathbf{i} > 0$ and $\mathbf{i}_s(t) = 0$ for $t > 0$ is one simple assumption, $\mathbf{i}_s(1) \equiv \mathbf{i}$ for all t another. The first assumes an exogenous source of infection that lasts only one year, the second a continuing exogenous source of infection. An intermediate and only slightly more complicated pattern would be to assume one value for the first year and a lower value for all subsequent years. We begin by assuming that $\mathbf{i}_s(t)$ has a constant value $\mathbf{i}_s(t)$ for all t .

The time-varying endogenous risk parameter

For the parameter $\mathbf{q}_s(t)$ we assume an initial level \mathbf{q} that continues until some time t_1 followed by a linear decline to zero at time t_2 . The initial value \mathbf{q} represents conditions in the population conducive to infection when the epidemic begins, which are assumed to remain unchanged until time t_1 . From time t_1 conditions in the population are assumed to change in a way that reduces the risk of infection for any given level of prevalence. These reductions in the risk of infection might result from behavior change, or from biological changes in the virus. No generality is lost in assuming that $\mathbf{q}_s(t)$ declines to zero, for any value above zero would be incorporated into the time invariant parameter \mathbf{k} . The assumption of piecewise linear change might of course be replaced by a more plausible curvilinear pattern. In present context, however, this would almost certainly have a negligible effect on the results, and linear formulation has the advantage of simplicity.

Mortality risks from HIV-related causes

We must also consider the level of risks of death from HIV-related causes for infected persons. These exert a powerful influence on the development of the epidemic because they control how long infected persons remain in the population after infection and therefore what level of prevalence results from any given cumulative incidence. Higher levels of prevalence feed back, in turn, to higher incidence levels.


Initial fittings of the model to the Uganda prevalence time series for 1985-1998 with the age-duration matrix \mathbf{a} of risks of death from HIV-related causes resulted in time trajectories of prevalence that rise and fall too slowly to give a good fit to the data. It is perfectly possible that this reflects problems with the data, which indicate that prevalence does fall extremely rapidly between 1992 and 1998. It is also possible, however, that the values of alpha estimate in the last chapter, which correspond to a median survival time of approximately 9 years, are too low, for they are derived from a single study (Masaka) representing only 90 incident cases. To allow for the possibility that the \mathbf{a} values derived in the preceding chapter are too low we allow for multiplying them by a constant “ \mathbf{a} level” parameter. Other things being equal, putting the value of this parameter greater than one increases these mortality risks, decreases prevalence, and results in a more rapid rise and fall of the epidemic.

Initial fits


The first fits to the Uganda prevalence data were made with the \mathbf{a} values presented in the last chapter, varying \mathbf{i} , \mathbf{k} , \mathbf{q} and the “start time” t_1 and “end time” t_2 . A plausible initial value for the start time was taken to be nine years, long enough for the epidemic to began to take sufficient toll to trigger preventive reactions. A plausible initial value for the stop time t_2 is less clear, but the initial value was taken to be 20 years. It was found, however, that no combination of \mathbf{i} , \mathbf{k} ,

and q would give a reasonable fit to the data unless this stop time was sharply reduced. Indeed, the best fit occurs when it is reduced to 10 years, implying a sudden reaction that reduces the (time-varying) component of endogenous risk to zero in a single year. This is completely implausible, obviously, but the result is worth showing because it illustrates the reasoning involved in applying the model to real data.

The following figure shows observed and fitted prevalence with $i = 0.001$, $k = 0.01$, $q = 1.28$, a start time of 9 years, and an end time of 10 years. The initial observation, for 1985 is assumed here to be for the eighth year of the epidemic, i.e., the epidemic is assumed to have started in 1976.




Though the model points fit the rise of the epidemic very well, they decline too slowly (if the data is believed), even though q remains positive for only a single year. Choosing a higher stop time, i.e., a slower preventive reaction to the epidemic, results in a model fit that goes much too high and fits much less well, as shown in the following figure.




The “rounded top” of the model prevalence trajectory reflects the relatively long period between infection and death. New infections enter the pool of infected persons and, as long as the \mathbf{a} values for risk of death from HIV-related causes are fixed at the levels estimated in the preceding chapter, remain there for an average of nearly 10 years no matter what happens to the values of the remaining parameters.

One interpretation of this result is that the Uganda prevalence data come down too fast between the peak in 1992 and 1998. An alternative interpretation is that the \mathbf{a} values estimated in the preceding chapter are too low. To test this interpretation we double the magnitude of these values, adjust the value of \mathbf{q} down to 1.0, increase the stop time to 14 years, and move the origin of the epidemic back by one year to 1975, resulting in the following fit.



The fit here is sufficiently close that there is no point in further manipulations of the parameters to improve it. The time over which reaction to the epidemic brings \mathbf{q} from its initial value of 1.0 down to zero is now five years, which might be regarded as minimal. The upward adjustment of the \mathbf{a} values is very considerable, but might be regarded as more plausible than a shorter reaction time.

A striking feature of this fit is the very low values of $\mathbf{i} = 0.001$ and $\mathbf{k} = 0.01$ values in relation to the initial value of $\mathbf{q} = 1$. Higher values of \mathbf{i} and \mathbf{k} would of course result in a higher level of endemic prevalence. One might speculate that the values of \mathbf{i} and \mathbf{q} required to generate the initial epidemic are inversely related. A very low value of \mathbf{i} will suffice if the initial value of \mathbf{q} is sufficiently large. A higher value of \mathbf{i} and a lower value of \mathbf{q} might accomplish the same result for the epidemic proper, though they would result in a higher endemic prevalence level. The speculation is easily tested. Doubling the value of \mathbf{i} to 0.002, for example, we find that reducing the value of \mathbf{q} to 0.85 gives a very similar fit, shown in the following figure.



Further experimentation shows that increasing i by a factor of 10, to 0.01, and reducing the initial value of q to 0.5 gives a fit that is noticeably poorer, rising too early and falling too slowly, and with an endemic prevalence level of well over 10 percent.

Discussion

Our objective in this application was relatively modest, to use the model to simulate an epidemic broadly similar to that observed in Uganda, as reflected in the available time series of prevalence data, and using all pertinent information on demographic and epidemiological parameters. The last condition entailed, in fact, a large proportion of the work involved, though this has laid the basis for investigations of the model beyond those given here.

Having accomplished the objective we now consider what has been learned and what lines of investigation might be pursued next.

The preceding section shows that, despite the many simplifying assumptions that have been invoked, the number of parameters is sufficiently large that many different combinations will provide a reasonable fit to the observed prevalence data. This is the rationale, in fact, for making many assumptions that might, *a priori*, be objected to on the grounds that they may be “unrealistic.” Relaxation of these assumptions and the large number of free parameters that would result would overwhelm initial attempts at fitting. Once experience has been gained with a more constrained model, refinements may well be appropriate.

It is evident even from this very brief exercise that fitting the model to epidemics in developing countries (and perhaps in many developed countries as well) will never be merely a matter of varying parameter estimates to obtain an optimum fit, but will involve judgements about the reasonableness of various parameter values. That the value of the time-varying endogenous parameter q should decline to one in a single year is unreasonable, even if this does give the best fit to the Uganda prevalence data.

When fitting perfectly specified models to perfect data it may be possible to draw conclusions by purely mechanical means. When fitting manageably complex models to real data, every discrepancy may be explained in at least two ways, as a defect of the model, or as a defect of the data. In the nature of the situation, there can be no formal, statistical decision rules in this context.

Conclusions can only be drawn on the basis of judgement informed by extensive knowledge of the model, of the data available in the particular application, and of experience with other models and other applications.

A powerful advantage of the model developed here, part of its conception from the very beginning, is the possibility of tracking demographic impact, over the long term as well as the short term, and with the kind of demographic detail that is normally incorporated in population projections. Once parameter values are decided it is possible to compute annual changes in crude birth, death and growth rates, expectation of life at birth, age pattern of deaths, and many other characteristics. A detailed demographic impact analysis will be appropriate, however, only when the model and its parameters are better understood than they can be from this preliminary application.

It is important, in particular, to gain a clear understanding of the way in which the various epidemiological parameters interact to produce both an epidemic and, possibly, sustained endemic prevalence. Some general questions may be asked of nearly any model. Do all the model parameters have independent effects, or is it possible to find a reparameterization that involves fewer parameters while generating the same range of outcomes? Is the model “ergodic,” i.e., is the eventual endemic level independent of the way in which the epidemic began and developed in the early stages? Other questions are particular to the general model developed here, e.g., what sort of differences result from putting the exogenous risk parameter \mathbf{i} to zero after the first year of the projection? Both kinds of questions may be usefully addressed through comparisons of the “bare bones” model with the full model as well as through analysis of the full model itself.

It is possible to explore questions of this kind independently of any particular set of empirical data, and to a limited extent this is desirable and even necessary. The range of possible questions and modeling experiments is so large, however, that analysis carried out independently of empirical applications risks practical irrelevance. Confronting the model with real data for actual epidemics is more likely to raise the right questions, “right” in the sense of useful for practical analysis and projection of particular epidemics.