

Comparing Uganda and South Africa

*Sexual Networks, Family Structure,
and Property*

South Africa and Uganda have been at the center of attention in the HIV/AIDS pandemic.¹ The first evidence of the HIV epidemic in Africa emerged in Uganda in the early 1980s.² It is the only African country that has shown an overall and sustained decrease in prevalence (until recently, when prevalence has apparently begun to rise).³ By contrast, South Africa was not forced to deal with the epidemic until the 1990s, and today has an HIV prevalence that is among the highest in the world. Since 1992, HIV trends in the two countries have moved strongly and consistently in opposite directions. This cannot be due simply to levels of sexual activity, since Uganda had a total fertility rate (TFR) of 6.9 in 2001⁴ and 7.1 in 2005,⁵ the highest in eastern and southern Africa. Uganda's TFR has remained roughly stable at least since 1955. South Africa's fertility, in contrast, has been relatively low and is declining,⁶ with a TFR of 2.7 in 2005,⁷ the lowest in Africa.

The explanation for the difference in the prevalence trends lies in differing *configurations of sexual networks* in the two countries. To demonstrate this, this chapter develops a mathematical model of prevalence trends and shows how this allows us to see differences in the topology of sexual networks. (A certain amount of mathematics is required here, but this is separated from the rest of the text and a reader who is not interested in the mathematics can skip over these sections.) Chapter 3 shows how sexual networks are embedded in and shaped by differing social contexts, especially with respect to marriage, reproduction, the

development of households, and the transfer of wealth and property. Together, these approaches help to explain the radical differences between HIV prevalence trends in Uganda and South Africa.

It is clear that in Africa, sexual networks are responsible for transmitting the virus in almost all adult cases of HIV infection. It is also increasingly clear that the configuration and dynamics of sexual networks—such as their periodicity or timing, their density, and their randomness or clustering—are as important as individual sexual behavior in understanding the dynamics of HIV transmission, incidence, and prevalence.⁸ For instance, it has been estimated that newly infected people with high levels of HIV in their blood are up to ten times more likely to infect others than are people with older infections.⁹ This estimate is based on research by Maria Wawer and her team on a sample of ten thousand people ages 15 to 49 living in forty-four villages near Uganda's border with Tanzania. If newly infected people with high levels of HIV in their blood have sexual contact with multiple partners during this time, they become extremely efficient transmitters of HIV in the overall sexual network.¹⁰ The efficiency with which HIV is transmitted in sexual networks is highly sensitive to its periodicity and to the number of links (sexual contacts) between participants.

The concept of *sexual networks*, as it is developed here, introduces a social dimension into research that has, for the most part, focused previously on the individual (knowledge and attitudes) and on behavior and practices. This involves a change in the scale of analysis: large-scale social formations as compared to small-scale individual behavior. It also involves a conceptual shift toward a social-epidemiological model that is able to accommodate society-level structures rather than simply statistical aggregates of individual behavior in so-called populations. If we take sexual relations to be *social relations* (not just behavior), and sexual networks to be types of *social structures*, we can see that HIV is transmitted at the population or national level by specific types of large-scale social structures, that is, sexual networks.

This approach constitutes a significant methodological departure from the standard statistical medical-epidemiological approach, and from those social-scientific approaches to HIV/AIDS that focus on the individual.¹¹ This can be seen by comparing the standard epidemiological model (shown in figure 3), produced by Rand Stoneburner and Daniel Low-Beer¹² with the model based on power-law curves shown in figure 5 (to which we turn later in this chapter). Stoneburner and Low-Beer use a Gaussian model (called a bell curve because of its shape) to

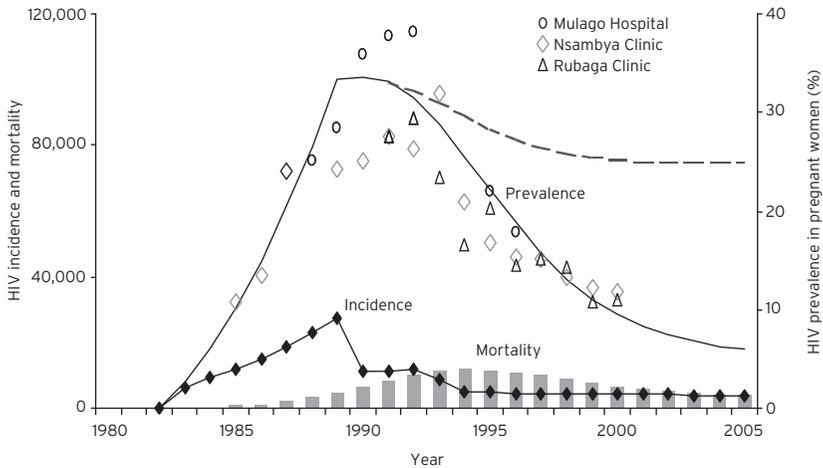


FIGURE 3. Stoneburner and Low-Ber's epidemiological model of HIV prevalence change in Uganda. (Based on Stoneburner and Low-Ber 2004: 717, fig. 3A)

approximate the empirical data. This model probably represents the best fit between the observed data (all from Kampala, drawn from the principal antenatal sentinel surveillance sites at Rubaga, Nsambya, and Mulago hospitals) and the standard epidemiological model.

Stoneburner and Low-Ber explain their model as follows:

Simulations of HIV incidence and prevalence in pregnant women, and mortality in the population aged 15 to 59 of Kampala, Uganda, from 1981 to 2005, showing simulated HIV prevalence under HIV incidence “intervention” (solid line) and “baseline” (broken line) scenarios compared with empirical HIV data from antenatal clinic sentinel surveillance sites in Nsambya (diamond shapes), Mulago (circles) and Rubaga (triangles). The intervention scenario, in which incidence rates are reduced by 80% over a 2-year period among 15- to 24-year-olds beginning in 1992–1993, produces declines in HIV prevalence consistent with empirical data. In the baseline scenario, incidence rates remained unchanged after 1990 and prevalence remains stable.¹³

Several problems are immediately apparent. In order to make the standard epidemiological model approximate a fit to the empirical data, the authors have had to assume that incidence (new infections) declined suddenly by 80 percent between 1989 and 1990.¹⁴ This is extremely unlikely and does not correspond to any of the recollections and stories about this period that Ugandans told us during research

conducted in 2003. The reduction in incidence shown in this model is only hypothetical, introduced in order to make the standard epidemiological model achieve a closer fit to the data. In fact, the decline in prevalence is such that even if incidence had been zero—extremely improbable—it still would not account for the data.¹⁵ Some researchers, however, have taken the hypothetical decline in incidence to be real. For instance, the authors of the influential 2002 U.S. Agency for International Development document *What happened in Uganda?* do this (repeating it again in 2006) when they argue that “it is most probable that HIV incidence in Uganda peaked sometime during the late 1980s.”¹⁶ In other words, to make the data (observed HIV prevalence) fit the model (a standard epidemiological curve) we have to assume that four out of five of all sexually active people suddenly stopped having unprotected sex for a couple of years.¹⁷ Without this sudden, dramatic, and improbable collapse in incidence, for which no confirming data exist, the dashed line wildly overestimates the actual HIV prevalence.

Indeed, *incidence actually increased annually* from 0.9 percent 1993 to 2.3 percent in 2003, even as prevalence was falling for those who tested at voluntary counseling and testing centers.¹⁸ The rise in incidence is consistent with a falling HIV prevalence only if incidence is rising in relatively small, isolated networks. If incidence were in fact rising across a randomly selected population that is widely linked in a common sexual network—as the standard epidemiological model assumes—then prevalence would also necessarily rise. This is not the case in Uganda. Instead, incidence is stable or rising in relatively isolated subnetworks (lately, especially among women and middle-aged people) within which HIV prevalence may rise quite rapidly, but outside of which prevalence may remain stable or fall. This is also consistent with the fact that Uganda has a very high TFR.

Indeed, since there are no population-based estimates of HIV in Uganda or hard evidence of behavior change before 1989, all must be conjectural.¹⁹ Even so, at the end of the Stoneburner and Low-Beer curve, the calculated hypothetical trend line (solid line) considerably *underestimates* HIV prevalence, which in reality levels off at a higher level than the standard epidemiological model would suggest. Furthermore, the curvature of the estimated (or fitted) curves tends in the opposite direction to the easily observable curvature of the actual data. The estimated curve in the period 1982–1992 is concave, or negative, but the actual data appear to trend in a convex way. In the period 1992–2002 the estimated curve shows a typical down-sloping logistic

or sigmoid form tending toward zero, whereas the actual data suggest a concave downward trend that levels off well above zero. It appears that the epidemiology model is simply of the wrong form (cf. figure 4 later in this chapter).

In Uganda, HIV remains endemic at higher prevalence levels than other virus infections even following the collapse in prevalence after the peak of an epidemic. This is because HIV does not confer immunity to infection on survivors (because there are none) as recovery from other viral infections does, and because the period of active viral infection (viremia) is much longer in the case of HIV (and generally bimodal or multimodal, rising at the beginning of infection and at the end). More important, however, the standard epidemiological model is based on the “normal” (Gaussian) probability curve,²⁰ which assumes that infection is randomly distributed in the population. Although this is roughly true for most infectious agents that are dispersed by water, air, and food supplies, it is not true of sexually transmitted pathogens, which rely on “wet” person-to-person genital contact. In other words, HIV is transmitted primarily through a social network of a special kind, the sexual network. This error has had much greater consequences, however, than the difficulty it presents to those who seek to develop mathematical models of the epidemic.

This problem in modeling HIV prevalence was noted in the early 1990s by two molecular biologists at the University of California, Berkeley.²¹ Peter Duesberg and David Rasnick argued that because HIV epidemiology did not fit the normal epidemiological patterns characteristic of almost all other viral and bacterial epidemics, it must not, therefore, be a viral epidemic at all. Basing their argument at first on data from the United States and Europe, they pointed out that “AIDS is highly non-random with regard to sex (86% male); sexual persuasion (over 60% homosexual); and age (85% are 25–49 years old). . . . From its beginning in 1980, the AIDS epidemic progressed non-exponentially, just like lifestyle diseases.”²²

They also pointed out that AIDS has no single defining disease²³ but is a syndrome, like obesity or illnesses caused by smoking, and that the vast majority of people who died of AIDS had a history of recreational drug use, or had taken antiretroviral (ARV) drugs. Since HIV infection in the population they had studied was also nonrandom, they concluded that HIV does not cause AIDS but was a product of toxicity arising from recreational drug use, lifestyle choices (especially ones that often accompanied the gay lifestyle in the United States and Europe in

the 1980s), and the drugs that were then used to treat HIV, especially AZT. If HIV were an infectious virus, they reasoned, it should infect people more or less randomly (that is, there would be a Gaussian distribution of infections), and should grow exponentially at first. HIV infections did not do this. In a 1996 publication, Duesberg and Ellison claimed—correctly at the time for U.S. data, but ignoring African data entirely—that HIV was still not heterosexually transmitted “outside of drug addicts” and that “by any measure, the war on AIDS has been a colossal failure.”²⁴

The core of Duesberg and Rasnick’s argument was the fact that the observed HIV trends in the data they examined did not conform to expected epidemiological patterns for viral pathogens such as flu virus or poliomyelitis.²⁵ Instead, it looked much more like the linear progression of illnesses caused by environmental factors such as pollution, illnesses caused by toxins, or lifestyle diseases.²⁶ Based on these and the other differences they observed, their conclusion that HIV was not a cause of AIDS seemed reasonable. Their conclusions were eventually shown to be wrong, but their critical observation about the shape of HIV prevalence trends was accurate.

Unfortunately, their findings reached the attention of the president of South Africa, Thabo Mbeki, a year or so later.²⁷ This was to have tragic consequences. Their logic convinced President Mbeki. In October 1999 he told the National Council of Provinces that AZT was poisonous,²⁸ and in early 2000 he set up the Presidential International Panel of Scientists on HIV/AIDS in Africa, which included Duesberg and Rasnick (see chapter 9). Then, in April 2000, in a strongly worded and very personal letter to United Nations secretary general Kofi Anan, U.S. president Bill Clinton, and British prime minister Tony Blair, Mbeki came close to saying that HIV did not cause AIDS, and urged world leaders to pay particular attention to the differences in the African HIV epidemic.²⁹ Mbeki’s position was drawn fully and directly from Duesberg’s work, although this was not directly acknowledged by the South African president. In a briefing paper for Mbeki that must have been received immediately before Mbeki penned his letter to world leaders, Duesberg stated the kernel of his argument that HIV does not cause AIDS, drawing on earlier work with Rasnick and Ellison. Duesberg wrote, “the African AIDS epidemic is not following the bell-shaped curve of an exponential rise and subsequent sharp drop with immunity, that are typical of infectious epidemics. Instead it drags on like a nutritionally or environmentally caused disease.”³⁰ Duesberg also

specifically warned Mbeki in his briefing document about those who opposed his views. Citing Mbeki's own letter to world leaders, which had been published by the *Washington Post* on 19 April, Duesberg wrote,

President Mbeki must also be warned about Dr. Joe Sonnabend's answer to the president's question about the epidemiological discrepancy between the "heterosexual" AIDS epidemic in Africa and the non-random, 85%-male epidemic in the U.S. (Mbeki's letter to U.S. President Clinton, *Washington Post*, April 19, 2000). According to Sonnabend's hypothesis, Africans acquire HIV heterosexually, because they simultaneously suffer from a long list of diseases, including "tuberculosis, malaria, other protozoan infections, bacterial diarrhoeal infections, pneumonia, plasmodium, Leishmania" etc. However, the very low AIDS risk of an African HIV-positive, compared to an American, calls this hypothesis into question. If the Sonnabend-hypothesis were correct, African HIV-positives should develop AIDS much more readily than their American counterparts. But the opposite is true. In fact according to Sonnabend most Africans should already have AIDS by the time they pick up HIV "heterosexually."³¹

This is remarkable logic by any measure, and wholly tendentious. But, by this time, Mbeki was fully and personally committed to what came to be called the dissident position. Duesberg and Rasnick had convinced him through their topsy-turvy logic (a) that the drugs used to treat AIDS were in fact the cause of it,³² and (b) that because of parasites, other bacterial and viral infections, and a generally higher disease load that Africans carried as a result of poverty, they probably already had AIDS by the time they contracted HIV!

Accordingly, in his address to the World AIDS Conference in Durban in April 2000, Mbeki declared that poverty was the real cause of AIDS. Predictably, this unleashed a storm of protest and dismay, but up to the time of this writing (early 2008) Mbeki has yet to renounce any of these beliefs. All of the dissident documents and his own dissident writings are still available on the Internet,³³ some of them through official African National Congress Web sites. The controversy, however, has succeeded in stopping him from expressing his views publicly. It has not prevented a rearguard action from the President's Office and the Department of Health aimed at delaying provision of ARV therapy and mother-to-child transmission protection.

The root of this tragedy lies originally in Duesberg and Rasnick's perfectly valid observation that the HIV prevalence curves are not typical of epidemics caused by other pathogens. They cited one example of an attempt to predict HIV prevalence that appeared in the *Journal*

of the *American Medical Association* in 1990.³⁴ This lack of fit to the “normal” curve is due to the fact that HIV is transmitted through sexual networks and is not randomly distributed. As we shall see, there are good reasons for this divergence from normal epidemic processes, reasons that do not mean that HIV does *not* cause AIDS, or that it is a disease of poverty or lifestyle.³⁵ It simply shows that, like environmental toxins and lifestyle diseases, HIV is spread through social structures (sexual networks) and is shaped by them.

Previous social research on HIV/AIDS has concentrated on the behavioral responses of individuals, especially their knowledge, attitudes, beliefs, and practices (KABP) with respect to sex and reproductive health. Work on sexual networks has previously focused on identifying risk factors and tracing contacts, and, in any case, is virtually absent in Africa. By contrast, the approach taken here directs focus away from individual behavior and risk in order to focus on the sexual network as a form of social structure, not simply contacts occurring as the result of behavior.

Unfortunately, while we have some data on sexual behavior in Africa, we mostly lack specific data on the details of sexual networks.³⁶ As early as 1992, Caldwell, Caldwell, and Orubuloye noted, “The AIDS epidemic in sub-Saharan Africa has revealed the inadequacy of our knowledge of the extent of sexual networking in the region.”³⁷ This was still true in 2006: empirical data on sexual networks in Africa, with a few exceptions,³⁸ do not exist. There are, however, ways to study the structure of networks indirectly in other ways.

A NEW APPROACH TO NETWORKS

Recent discoveries—many of them since the late 1990s—in computer science and mathematics about the behavior of networks provide new conceptual tools that can help us to illuminate the hidden processes in sexual networks.³⁹ Researchers have examined large-scale complex networks of many kinds, including the Internet (the physical infrastructure), the World Wide Web (the hyperlinks that make the Internet work), networks of neurons, food webs in small ecosystems, secret communication networks among terrorists, and citation networks among physicists. They have found that there are a number of invariant rules that govern how networks are configured and, more important, how they change and evolve. For instance, research has shown that these and other networks, irrespective of their size, have a “small world”

configuration in which all members of the network are linked with, on average, as few as two or three links between any two nodes in the network, and at most six links—known as “six degrees of separation” in social networks.⁴⁰

These findings can be usefully applied to the study of sexual networks. Analysis of how HIV prevalence changes over time—the trend—can provide information about the gross structure of sexual networks for which we lack detailed empirical data. This chapter uses these new models to examine the shapes of the trend lines based on least-square approximations to a power-law model of changing HIV prevalence in order to show that the shape of these trends suggests differences in the structure of sexual networks in Uganda and South Africa. Findings from trend-line analysis are supported by analysis of differences in social structure that seem most likely to have an impact on the configuration of sexual networks, including patterns of kinship, marriage, household structure, inheritance, and wealth.

COMPARING HIV TRENDS IN UGANDA AND SOUTH AFRICA

HIV prevalence in Uganda rose very rapidly in the 1980s and early 1990s until 1992, when it began to decline very rapidly.⁴¹ When it first began to be measured in 1985, it was already above 10 percent at antenatal clinics. It is likely that HIV had been endemic at a low level in Uganda for at least a decade, and possibly for generations, before it was first tentatively measured in 1982.⁴² By 2002, the decline in HIV prevalence had stabilized at around 7 percent, fell slightly, and then possibly crept up again by 2007.⁴³ The direction of the trend after 2002 is still not clear.⁴⁴ On the other hand, South Africa has seen a steady rise in the prevalence from the beginning of the 1980s until the present, with as yet little evidence of a leveling off.⁴⁵ In 1992, the overall prevalence in Uganda was around 24 percent, with rates as high as 29 percent in Kampala and averaging 20 percent in the four other major towns in Uganda.⁴⁶ In 2005, the overall prevalence in South Africa nationally was similar to that in Kampala, Uganda, thirteen years earlier. Prevalence in some age and sex groups and in some regions in southern Africa now exceeds 30 percent, a level never reached in Uganda.

The vast difference between the two countries is illustrated most starkly in figure 4. This simply shows the temporal relationship of changing trends in the two countries. The data on which these trends are based are, of course, much messier. There is also a large margin of error associ-

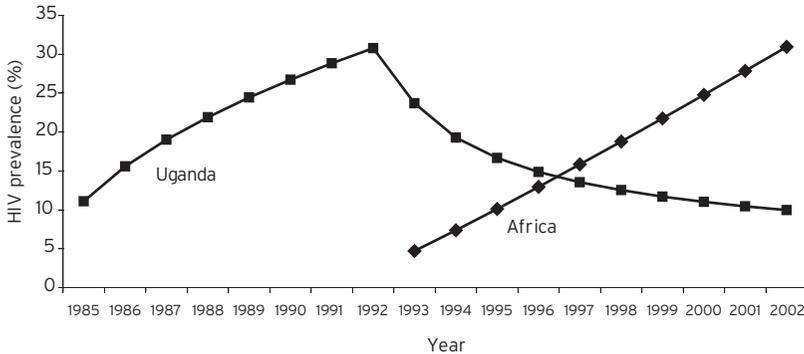


FIGURE 4. Comparison of HIV trends for Uganda and South Africa, 1985–2002. Prevalence trends in Uganda and South Africa are starkly different.

ated with each data point. Since we are concerned only with longer-term trends, however, the margin of error associated with each point can be eliminated as a concern.⁴⁷ The fact that these trends are highly consistent with each other for different data sets within each country also increases the reliability of these findings. Since we are trying, first of all, to compare two countries in the broadest possible strokes, however, elimination of the noise in the data allows us to see the overall relationships in time and in the shape of these trends. In the following analysis, it is necessary to separate out different periods for close analysis.

We first compare these trends for the eleven-year period (1992–2002 inclusive) when South African and Ugandan trends were moving in opposite directions. When we plot these two different national trends against each other, several significant patterns emerge during that decade.⁴⁸ Figure 5 shows the trend for South Africa as a whole against two sets of data for Uganda.⁴⁹ The triangles represent aggregated data for Kampala (drawn from the same antenatal HIV surveillance sites used in Stoneburner and Low-Beer’s model, shown earlier in figure 3). In this plot, median values for these three sites have been used, because each set of data alone shows a broadly similar trend. Squares represent the data from Uganda’s four largest cities with antenatal clinic sentinel sites: Mbarara in the southwest, Jinja in the south central area, and Tororo and Mbale in the east. Diamond-shaped markers represent data points for antenatal clinic sentinel sites in South Africa.

For each set of data, a power, or power-law, trend line is calculated and plotted (see explanation boxes). Curves of this sort are associ-

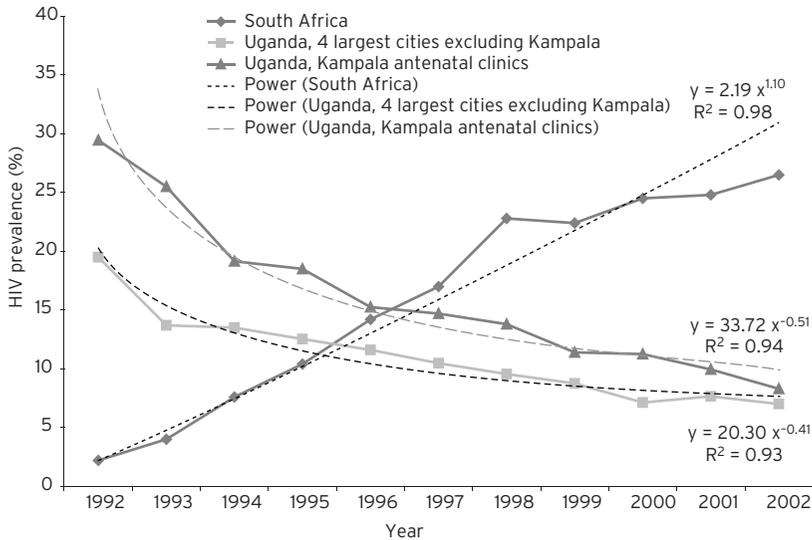


FIGURE 5. Comparison of HIV falling and rising power-law trend lines in Uganda and South Africa, 1992–2002.

ated with complex network systems.⁵⁰ The remarkable fact that the Ugandan data follow a power-law trend much more closely than it does the standard epidemiological model, shown in figure 3, suggests that efforts to understand the HIV epidemic in the same terms we use to model and visualize other epidemics is likely to fail. Something else is at work here. Since power-law trends are closely associated with networks of all kinds, it suggests that we might gain insight by examining the social networks that transmit HIV. Power-law trends are also associated with the sudden changes in chemical and physical systems called phase changes, such as when water (a liquid) suddenly turns to steam (a gas) at the boiling point, or when it turns to ice (a solid) at the freezing point. It turns out that phase changes are analogous to the changes in HIV prevalence in Uganda.

Where sexual networks are infected with HIV, the prevalence of infection will be a function of the number of sexual contacts (links) between people (network nodes) in a network of a specific configuration, and, over time, may behave in different ways depending on how the network is structured. For Uganda, it changes in a nonlinear way that is characteristic of highly structured complex networks with hubs and/or clusters and of networks that exhibit fractal structures. This is distinctly different from a linear growth that would be seen in simple

THE MATHEMATICS OF THE TREND LINE

The trend line (heavy dark line imposed over the data points) shows the closest fit of the data to a curve described by a power-law equation of the form $y = ct^a$. Here y is *HIV prevalence* expressed as percentage. The t is *time* in years, for this is how the available data are given to us. The symbol c is a proportionality constant that simply adjusts the scale of the graph, and a is the calculated exponent that allows the trend to fit the data.

The R^2 value estimates the closeness of fit, with values approximating 1 indicating a close fit. Here, the trend lines fit the data very closely, with values approximating 1.

When $a = 0$, the curve is simply a horizontal line in which $y = c$, a constant value. When $a = 1$, we have a simple linear relation, that is, a straight line. The South African data show an exponent close to 1, in which $y = ct$.

Fractional exponents where the absolute value of a lies between zero and one, $0 < |a| < 1$, however, yield curves like those seen in the Uganda case. If the exponent is negative, the curve has a concave shape with what is called a fat tail: the prevalence, y , declines rapidly at first and then tails off very slowly without reaching zero, as in figure 5. Positive values of the exponent yield the inverse "growth" pattern, shown by the trend line from 1985 to 1992 in figure 6.

This plot uses Microsoft Excel's built-in least-squares fit to the power-law model.

diffusion models of transmission, or exponential growth of unbounded biological systems, or the sigmoid curves of growing populations ultimately constrained by the carrying capacity or saturation of their environments. Specifically, the Ugandan curves lie halfway between stability ($a = 0$) and linear growth ($a > 1$) that would be characteristic of simple diffusion. This behavior is the signature of complex network organization that is fractal, or similar at different scales.⁵¹

The two trends for Uganda—one for Kampala, the capital of Uganda (comprising median values for antenatal clinics at Rubaga, Nsambya, and Mulago hospitals), and one for the other four major cities for which we have good data—both show approximately equal declines during this period.⁵² Both curves conform closely to each other, with the curve for Kampala reflecting higher HIV prevalence in the city than in the smaller regional cities. Trend curves for rural data and smaller towns (Moyo, Mutolere, Masindi, Hoima, Kilembe, Pallisa, Aber, Lwala, Soroti, Matany, Kagadi, Arua, Lacor, and Nebbi) are similar in shape, but lower in magnitude. Data for these sites are incomplete and therefore not presented here. The trends drop steeply for the first several years from 1992 to 1995, and then begin to level out toward an overall prevalence of around 6–7 percent. This is confirmed by the prelimi-

nary results of a national survey of HIV prevalence in Uganda based on voluntary blood tests conducted on a representative statistical sample for the entire population that also shows an overall prevalence of about 7 percent.⁵³ Equally significant is the fact that HIV levels stabilize but do not tend toward zero as ordinary viral epidemics would. This fact is accurately modeled by the power law, but not by the “normal” (Gaussian) distribution curve.

The decline follows a period from the early 1980s to 1992 in which HIV prevalence rose rapidly in Uganda toward a peak in 1992. The transition from rapid escalation of HIV rates in Uganda from the early 1980s up to 1992 was as remarkable as their decline. The period of increasing prevalence is also closely modeled by a power law. The transition between increasing HIV prevalence and decrease is very rapid (see figure 6 later in this chapter).

By separating the period of increasing prevalence (1985–1992) from the period of decreasing prevalence (1992–2002), we can calculate two discontinuous trend lines that show similar but inverted patterns.

The fact that these two trends were inverses of each other, with closely similar absolute values of the power-law exponent, suggests that whatever was happening in the overall sexual network to cause an increase in HIV was also happening during the period of decreasing HIV prevalence, but in a way that was somehow inverted. Moreover, this trend-line analysis suggests that the changes in HIV prevalence were functions of the configuration of the overall sexual network, rather than simply statistical changes in human sexual behavior. In other words, we do not have to assume that 80 percent of the sexually active population stopped having sex between 1991 and 1992, as Stoneburner and Low-Beer are driven to assume, given the nature of their model. As we shall see, it is very likely that what was happening during the period of increasing HIV prevalence in Uganda was an increase in the number of links (sexual contacts) between persons or groups of people (such as within villages, around nightclubs or bars, or at funerals or other ritual events), and especially an increase in the links between highly infected persons or clusters that were highly linked to other persons or clusters, such as soldiers, transportation workers, “sugar daddies/mamas,” and others. During the period of decreasing prevalence, overall connectivity decreased, but especially the links between highly infected persons or groups in the network and the general population decreased. This would explain the power-law shape of the trends since these are the signature of networks connected in this way.

POWER LAW: THE MATHEMATICAL DETAILS

Unlike an *exponential* prevalence curve, the *power-law* curve is simply a function of time, t , raised to a constant power. The exponent is a real number; that is, it may be a positive or negative whole number or fractional number. By contrast, the exponential curve expresses a situation in which the prevalence is a function of itself: that is, the higher the prevalence, the faster the rate of increase, or the lower the prevalence, the lower the rate of increase, usually (in real-life systems) up to some limit imposed by the availability of resources.

The trend of the decline in Uganda conforms to a power-law decay from the peak reached in 1992. The decline in HIV prevalence (P_{HIV}) as a function of time (t , in years, in accordance with the data sets) is estimated by the expression

$$P_{HIV} = 0.3372t^{-0.5087}$$

for Kampala antenatal center test sites, and by

$$P_{HIV} = 0.203t^{-0.4129}$$

for the four other large cities (where t is time in years, in accordance with the data). The two trends are very similar, and both have a negative noninteger exponent a , $0 < |a| < 1$. The trend for the period of *increasing* HIV prevalence, 1985-1992, for Kampala (figure 6) is estimated by

$$P_{HIV} = 11.09t^{0.4912}.$$

All power-law exponents (-0.5087, and -0.4129 for the period of decreasing HIV prevalence, and 0.4912 for the period of increasing prevalence) are approximately of the same magnitude, $|a| \approx 0.5$. The similarity in these values suggests that we are dealing with the same form of network. In the first period (1985-1992), links were rapidly forming, making the sexual network an increasingly *efficient* transmitter of HIV. In the second period, links were rapidly breaking, resulting in a *decreasing efficiency* of HIV transmission in the sexual network.

Previous work on networks has assumed that they are static and can be represented adequately by graphs (drawings) that represent the links between nodes at some arbitrary time or over some arbitrary period of time.⁵⁴ This is not adequate for our purposes, however. Sexual networks are dynamic, not static: over time, they grow or shrink, their configuration may change, and the number of links in them may increase or decrease. We need a model of networks that takes this into account.

Fortunately, the power law describes just such changes in a network. Research on the evolution of the hard-wired Internet and software-based World Wide Web (WWW) has shown that “the networks are becoming *denser* over time, with the average degree [the average number of links] increasing (and hence with the number of [links] growing super-linearly in the number of nodes). Moreover, the densifi-

THE NETWORK POWER LAW

Leskovec, Kleinberg, and Faloutsos (2005) have found that “as the graphs [of networks such as the Internet and World Wide Web] evolve over time, they follow a version of the relation”:

$$l(t) \approx n(t)^a$$

where $l(t)$ and $n(t)$ represent the number of links (l) and the number of nodes (n) at some time t , and a is the power-law exponent. The expression $l(t)$ simply states that the number of links is a function of time. This is a power-law relation that means that the number of links in these networks increases in proportion to the number of nodes raised to some power represented by the exponent a . This expression describes the growth of a network. Leskovec, Kleinberg, and Faloutsos (2005) call it the “densification power law” or the “growth power law.” Since their research concerns the Internet (the physical electronic infrastructure that connects computers) and the World Wide Web (the software links connecting users that are followed by browsers such as Netscape, Firefox, and Internet Explorer), both of which are growing rapidly, they do not consider the case in which networks shrink by losing links. This is what happens when HIV prevalence declines. In this case, the inverse of the “growth power law” — in other words, a “decay power law” — exists when the exponent is negative, as follows:

$$l(t) \approx n(t)^{-a}$$

This relation says that the number of links between nodes decreases more rapidly (sublinearly) than the number of nodes (persons). This relation closely models the period of decreasing prevalence in Uganda.

HIV prevalence in the population, then, will be proportional to the number of links (sexual contacts) that connect the nodes (persons) in the sexual network as a whole, itself a subset of the entire population. Thus, HIV prevalence will vary according to the number of links, which, in turn, is proportional to the number of nodes raised to the power a , the “densification exponent” that is characteristic of all such complex preferential attachment networks. Accordingly, a power law that describes HIV prevalence trends emerges:

$$P_{HIV} \approx l(t) \approx n(t)^a$$

cation follows a power-law pattern.”⁵⁵ What this means is that empirical observation of very large networks such as the Internet and WWW shows that the number of links between nodes increases more rapidly (superlinearly) than the number of nodes. This is caused by what has been called preferential attachment, or a “rich get richer” phenomenon.⁵⁶ According to this, the establishment of links to new nodes tends to favor those that are already highly linked. Nodes that are already popular, such as Yahoo.com or AOL.com, attract more links than other Web sites do. This is also true of neurons in simple organisms, networks of Hollywood actors, or famous scholars whose works are

most often cited. According to Barabási, these patterns, “potentially present in most networks, could explain the power laws we spotted on the Web and in Hollywood,” among many others.⁵⁷ In the case of sexual networks that transmit HIV, clusters of intensive sexual contact produce higher likelihoods of HIV transmission. In other words, with respect to HIV, the “rich” (highly infected clusters) get “richer” (that is, prevalence increases more rapidly than it does in sexually less well-connected clusters or categories of people).

THE POWER-LAW MODEL AND HIV TRENDS

It is intuitively obvious that HIV spreads through links in the sexual network. Thus, nearly all seropositive people will be members of the sexual network, in which some people overall, and *all* seropositive people, necessarily have direct or indirect sexual contact with more than one other person. In fact, the set of HIV seropositive people is a subset of the network of all sexually active people in the larger population. Infected individuals constitute a sample or subset of the set of all members of the sexual network that has been “selected” by HIV infection. Within the network, all HIV-positive people are necessarily linked by some set of links to all other HIV-positive people, and to some HIV-negative people. Some members of the set of HIV-positive people are likely to have many sexual contacts with others, especially during periods in which HIV prevalence is increasing. They may constitute hubs in the network, and are efficient transmitters of HIV. In fact, there may be clusters of HIV-positive people whose links with one another and with some HIV-negative people are quite dense. These persons or clusters can act as transmission centers or hubs. Sexual networks, however, usually have long filaments of contacts with cross-links to other filaments and to hubs. These filaments can act as transmission lines in the way electrical power lines do. The efficiency with which HIV is transmitted overall, then, will depend on the specific configuration of the network. This is what we find in the data for Uganda. The overall picture of increasing *and* decreasing HIV prevalence for Kampala data is shown in figure 6.

It is important to note that the pattern observed here does not look like the patterns seen in other types of epidemics caused by “normal” bacteria or viruses. The composite curve in figure 6, which fits the data quite closely, shows a remarkable and sudden reversal. Such trends suggest that the dynamics of the system are unstable, and that something

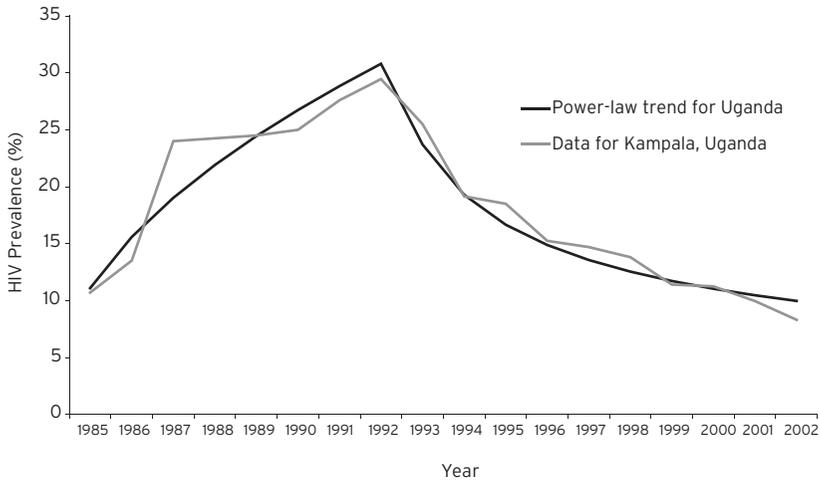


FIGURE 6. Trend for changing HIV prevalence in Kampala, Uganda, 1985–2002.

like a phase transition or catastrophe event occurred in 1992. A sudden shift of this sort—a tipping point⁵⁸—is most likely to be caused by a rapid reorganization of the overall structure of the sexual network that transmits HIV rather than by incremental and local changes in HIV incidence caused by behavioral change, morbidity, or death alone. Much recent work in the mathematical analysis of networks shows that nonrandom networks (such as the Internet, but also social networks, scholarly citations, networks of neurons, and food chains in ecological systems) all show power-law patterns. Steven Strogatz, one of the leading innovators in this field, remarks, “Power laws hint that a system may be organizing itself. They arise at phase transitions, when a system is poised at the brink, teetering between order and chaos. They arise in fractals, when an arbitrarily small piece of a complex shape is a microcosm of the whole.”⁵⁹

This analysis of the Ugandan HIV data strongly suggests that just such a phase transition occurred. It also suggests that the structure of the network is fractal, that is, it exhibits a similar structure at different scales of analysis. This would be consistent with a highly clustered network in which different clusters of sexually active and therefore interlinked individuals are further grouped into larger clusters, and so on. In such a network, each cluster has fewer links to other clusters, both within the same order of magnitude (sections of a village) and at higher orders of magnitude (regions, language groups, age groups)

WHY A POWER LAW LOOKS LIKE A STRAIGHT LINE IN THE CASE OF SOUTH AFRICA

The trend line equation for the South African data is approximated by the power-law equation

$$P_{HIV} = 0.0219t^{1.1048}$$

where the exponent is close to 1.

As we have seen, where the exponent in the power law is 1 or close to 1, the graph of the equation approximates a straight line because any value raised to the power of 1 is simply that value:

$$f(t) = t^1 = t.$$

In the South African case, this corresponds to a more or less straight line sloping upward over time, meaning that the increase in prevalence is simply directly proportional to time, with a small multiplication factor of 0.0219 that increases the slope ever so slightly upward.

than it does within itself. This structure permits what amounts to the phase transition that we see in figure 6, in which the network suddenly became much more resistant to transmission of HIV virus around 1992. *This would be consistent with a model in which a highly clustered network with limited links between clusters suddenly lost those interlinking long-distance connections, resulting in a phase transition of the network.* Through a process similar to water freezing and suddenly becoming “lumpy” with ice crystals, and then becoming hard as it freezes to a solid, it seems that beginning about 1992 the Ugandan sexual network returned to a lumpy, clustered state. In other words, while people continued to have sex, possibly with multiple partners, persons or clusters that acted as hubs in the network largely disappeared. This is consistent with the data on changes in sexual behavior in Uganda.

By contrast, the trend for South Africa (figure 5) is approximated by a straight line. This would not immediately suggest that networks were involved, except for the fact that we know they are. What it does suggest is that the South African sexual network permits a more or less unimpeded flow of HIV through it, roughly like liquid paraffin diffusing through a wick in a lamp, or a strong smell diffusing through a room from its source. It takes a bit of time, but it moves at a constant rate. It would seem that something like this was happening in the South African sexual network.

NORMAL EXPONENTIAL GROWTH OF EPIDEMICS

In their early stages, epidemics normally grow exponentially, that is, the rate of change in prevalence of some infectious agent (for example, a flu virus) in the population is at first proportional to the number of persons infected: the greater the number of people already infected, the greater the rate at which more people will become infected. This relation can be expressed mathematically by the ordinary differential equation

$$\frac{dy}{dt} \approx y$$

meaning that the change in y with respect to time, t , is proportional to y . This has as its solution in y the simple exponential growth equation,

$$y = ce^{kt},$$

where c is the initial value (in this case 1), k is the proportionality constant, and t is time represented on the x -axis. The e , a constant value, is a special type of number, called a transcendental number, like π ; $e = 2.7183 \dots$. It can be thought of as the limit on the returns from compounding interest (versus simple interest) as the number of periods of compounding interest are increased. It is thus a natural value for the base of the exponential function.

The South African data show an approximately linear progression during the period 1992–2002. A linear progression would mean, in this case, that HIV prevalence (the y -axis) is simply a function of time (the x -axis), or, in other words, that prevalence is growing steadily with time and that the *rate* of change was not itself changing. We must keep in mind that the data do not preserve spatial characteristics, so what they show is a gradual trend toward saturation of the *entire sexual network* with HIV infection. This in turn suggests that the sexual network is relatively homogeneous and randomized, that is, that all members are likely to have contact with people both close and far away, or within some category (linguistic, cultural, age) as well as with others outside this category. It has been shown that people in multilingual communities in Uganda are far more likely to have extramarital sexual contact than those within monolingual communities.⁶⁰ Since virtually all South Africans live in multilingual communities and almost all speak several South African languages, this may predispose them toward greater sexual contact across the network. It also suggests that most sexually active people have multiple partners, and that most have more or less the same number of

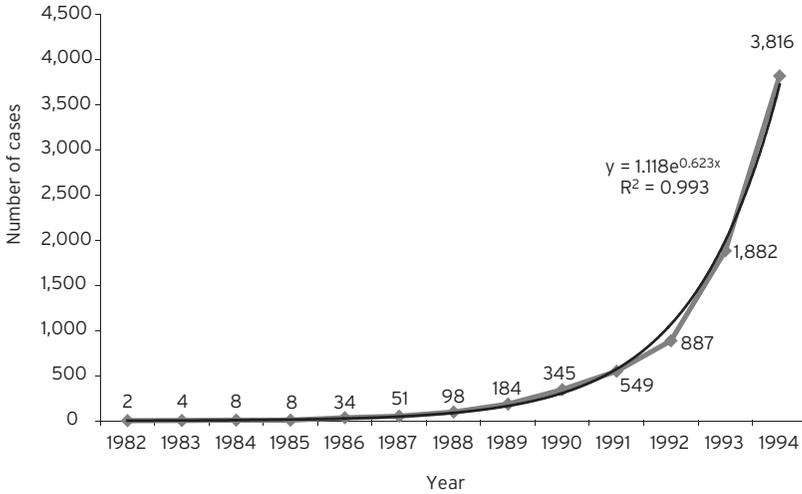


FIGURE 7. The earliest period of growing HIV prevalence in South Africa, 1982–1994. (Data source: United Nations Programme on HIV/AIDS 2002)

partners (links) in the sexual network. The behavior of the trend during this period suggests a nearly free diffusion of HIV through the population, as if the growth of the epidemic is a function of time alone, rather than—what is usually the case in epidemics—prevalence over time being a function of the number of infected people already in the population. Furthermore, it suggests that the system (all sexually active people linked in the HIV transmission network) is not yet close to saturation or phase transition.

But the growth of HIV prevalence in South Africa did not suddenly begin as a linear increase. In its earliest stages, it grew exponentially, as it did in the earliest periods of infection in the United States and other places for which we have relatively complete data. (We do not know what happened in Uganda at first, because HIV and AIDS were unknown and not tracked before HIV had already infected up to 10 percent of the population in certain parts of the country.)

Further analysis of South African data shows that the earliest period of the HIV epidemic is most accurately described as *exponential*. In its earliest stages, HIV prevalence increased freely, as if it encountered no barriers, that is, as if it were like a virus that could be spread by and infect anyone. This could be true only if the sexual networks in which

it was propagated were multiply and randomly connected. Viruses and bacteria that are free to move easily, and where there is no previous immunity, grow exponentially in a population before they either evoke immunity in previously infected people, or die off as their hosts succumb to lethal pathogens (in which case, they decline exponentially to zero or near-zero levels).

Exponential curves generally emerge in natural systems where growth is unrestricted, such as the growth of any population (dividing cells or free growth of viruses, bacteria, or other organisms) that is not limited by resources (food, light, or space). Such trends occur where the rate at which prevalence changes is a function of prevalence itself, that is, the more infections that already exist, the greater will be the increase, up to some limit such as death, recovery from infection, or lack of vulnerable individuals. The earliest South African data show precisely such a trend (the R^2 value of 0.9935 is a virtually perfect fit to the exponential trend line), as illustrated in figure 7. This trend, taken together with the later data, describes an almost perfect growth curve in which limits to growth have not been reached. It also strongly suggests that HIV transmission in South Africa occurs in a fully connected and randomized network without internal structural limits such as clustering or formation of core groups of endemic but limited infection.

(Re)Organization of Sexual Networks and Prevalence Change

The patterns we see in prevalence trends strongly suggest that sexual networks in Uganda are highly clustered or lumpy, with many locally dense subnetworks that have fewer connections between them. In South Africa, networks appear to be randomized and highly interlinked. The catastrophe event or phase transition in the Uganda prevalence trend resulted from a fairly sudden reorganization of the sexual network, during which the density of links between local high-prevalence, high-density networks—clusters or lumps in the network—and the rest of the population were suddenly reduced. In South Africa, this has not happened because the network is much more uniform, with large numbers of people multiply connected to one another.

In Uganda, the centers of infection were either clusters of highly interlinked and infected persons around a bar or night club, for instance, or else highly sexually active persons such as soldiers or truck drivers. People involved in culturally sanctioned networks of sexual exchange

(e.g., wife-sharing, or sexual cleansing after the death of a spouse) are also likely hubs. In terms of this model, it does not matter whether these links were severed by behavior change (using a condom or reducing the number and/or frequency of sexual contacts) or by death or other factors. The result was the same: fewer links between highly infected hubs, clusters or filaments and the rest of the population. Since infected and highly sexually active hubs (whether individuals or groups/clusters) are highly efficient in spreading HIV, reducing links from these hubs creates a nonlinear decline in the transmission of HIV. If some percentage reduction in sexual links occurs (say, r percent), then this will have little effect within densely connected subnetworks, but it will have a significant effect on linkage between clusters. In fact, r percent reduction of cross-cluster links will effectively isolate some parts of the population from infection. The effect overall will be a reorganization of the total sexual network in which prevalence will be nonlinearly related to r . This can be described as a phase transition. In other words, the rate of change in HIV prevalence was not directly proportional (linear) to the sum of changes in individual sexual behavior or death but was instead a function of the change in the organization of the overall network, which had the effect of accelerating the decline in HIV prevalence.

The structure of the South African sexual network appears, from this analysis, to be quite different. In particular, it seems that the sexual network spans the nation (and probably also much of the southern African region, including, at least, Botswana, Lesotho, and Swaziland, which share languages, culture, and population freely with South Africa). There appear to be few clusters or hubs in a relatively homogeneous network of sexual relations that ultimately incorporates most sexually active people (and certainly all those with more than one partner). In such a network, HIV spreads according to normal growth curves (logistic or exponential) and has so far not reached a natural limit. In fact, we cannot, as yet, say what that natural limit might be.

Given the lack of actual empirical data on sexual networks, then, it appears that analysis of trends may serve as a proxy measure for the structure of sexual networks. Since we know that the configuration of sexual networks is critical to the transmission of sexually transmitted infections, it follows that such knowledge can be crucial in understanding the progress of the epidemic. This, in turn, can help us to create and implement interventions that might be more effective than what we have seen so far in Africa.

Perhaps the most important lesson that we may derive from this new understanding of networks and HIV transmission is that significant change in HIV prevalence is a property of the social network rather than of individual behavior. We must redirect our attention, then, from the scale of the individual—behavior, psychology, risk assessment, and so on—to the scale of the social. This represents a radical shift in perspective.