The Impact of Measurement Error on Survey Estimates of Concurrency

Martina Morris
Departments of Sociology and Statistics
Pennsylvania State University, USA

John O'Gorman
Department of Statistics
Pennsylvania State University, USA

IUSSP Conference on Measurement of Risk and Modeling the Spread of AIDS

Copenhagen, June 2-4 1998

The authors would like to acknowledge the help of Mirjam Kretzschmar and Mark Handcock, who provided insightful suggestions and programming assistance at key stages of this project.

Background

In the absence of a vaccine for HIV or a cure for AIDS, behavioral change remains the only method of prevention. For sexually transmitted HIV, this means changes in the patterns of sexual behavior. While condom use and reductions in rate of partner acquisition have been the most obvious targets for behavioral intervention, a handful of studies have shown the importance of other factors, in particular, the level of concurrent partnerships (Kretzschmar, 1996; Morris, 1997). Concurrency, i.e., partnerships that overlap in time, prevents the pathogen from becoming locked in a partnership, and creates a long, loosely connected web of relationships in the population. In simulation studies, concurrency has been shown to have very large effects on the spread of a pathogen through a population, raising the growth rate of the epidemic linearly and the prevalence exponentially, even in populations with low rates of partner acquisition. Such a pattern, with relative low numbers of lifetime partners but high levels of concurrency, has been shown to be a good description of sexual behavior in, for example, Uganda.

The difference between partner acquisition rates and levels of concurrency is not simply an academic issue. In populations where the average number of lifetime partners is low but concurrency is common, public health messages that stress the importance of having fewer partners may be ineffective – most persons already have few partners, and are unlikely to perceive themselves at risk. For this kind of population, a message stressing the importance of monogamy is much more important. This suggests that an effective public health campaign must have accurate information to understand the needs of the specific target population.

Methods for collecting data on concurrency have been developed, and these typically rely on the reporting of dates. Little is known about the accuracy or reliability of the dates collected in this kind of survey setting. In this paper, we use simulation to examine the impact that date reporting error may have on estimates of concurrency from survey data. We focus on two types of reporting error: unit heaping, which is often imposed by the survey instrument, and recall error.

Methods

The survey instrument used for the observed data

To obtain information from a respondent about whether their sexual partnerships have (or do) overlap, there are two basic approaches: direct and indirect. The direct approach would be to ask the respondent a question like, "Did you have sex with any other persons while you were sexually involved with your most recent partner." There are several difficulties with this approach, ranging from lack of precision to increased risk of non-disclosure. In most cultures, having more than one sexual partner at a time is regarded as inappropriate, at least for women, and direct requests for this information are likely to meet with socially appropriate responses, regardless of the truth. The alternative survey approach is to use a local network questionnaire module. Under this approach, the interviewer asks a block of questions about the respondent's most recent sexual partner, then repeats the block for the next most recent partner, and so on until time or recall limits are reached (typically 3-5 partners, the "local network"). To establish concurrency, the module will include a question on the date the respondent first had sex with that partner, and the date they last had sex with that partner. At the stage of data analysis, these dates can be used to construct the relationship interval for each partner, and the intervals can be checked for time overlap. The respondent is never directly asked to explicitly divulge the information, and if these questions are interspersed with others in the partner block, they may occasion little embarrassment or discomfort.[1] The resulting information can be depicted graphically by Figure 1. In this example, the respondent reports 3 partners, with partners 2 and 3 concurrent for about 1 year.

This approach provides two opportunities for measurement error in defining the interval for each partnership, three opportunities in defining concurrency (one only needs to establish whether the end date of the prior relationship falls in the most recent interval), and four opportunities in calculating the length of the overlap. We will use this scenario to model the

[1] While the event history calendar often used in demographic settings could be proposed as a third alternative, it also requires the respondent to explicitly identify the overlap in sexual partnerships, and increases the risk of non-disclosure.

possible impact of date reporting error on the existence of concurrency, and the overlap interval, and the partners that are concurrent in a simulated survey setting.

Simulation Setup

The simulation involved four steps: (1) exploratory analysis of an actual data set to set the targets for the "true" values of the dates; (2) construction of a simulated "true" data set with dates reported in days; (3) addition of the heaping error induced by choice of reporting units (days, weeks, months, years ago); and (4) addition of recall errors. The final output was two datasets: the "true" dates, and the dates with reporting error.

For exploratory analysis to set the target values, we used a sexual network study from Uganda, the *Ugandan Sexual Network/Behavior Study for HIV Prevention* (Wawer, 1993). This is a population-based survey of about 1600 persons aged 15-45 in Rakai District, Uganda, fielded in 1993/4. The survey instrument contained a local network module with blocks of about 75 questions for each of the respondent's three most recent sexual partners (or fewer if the respondent had had less than three partners). We used these data to provide estimates for the following parameters:

1. the total number of spousal and non-spousal partnerships reported by the respondent
2. the rules for reporting units heaping into weeks, months, or years.
3. the distribution of partnership lengths
4. the distribution of sexual "age" (years since the respondent became sexually active)
5. the frequency and distribution of concurrency

These estimates were used as target values when constructing the simulated "true" data set. The target estimates are shown in Table A2 in the appendix.

The simulation can be thought of as follows: a person can be in one of the following four states: 0 ,1, 2, or 3 partners. The transition rates for gaining and losing a partner are assumed to be constant, conditional on the current partnership state:

Parameter	Description
$\square_1$	P(gain one partner | currently 0 partners)
$\square_2$	P(gain one partner | currently 1 partner)
$\square_3$	P(gain one partner | currently 2 partners)
$\square_1$	P(lose a spousal partner in a given interval)
$\square_2$	P(lose a non-spousal partner in a given interval)
$\square$	P(partner is spousal | partner is gained)

The parameters $\square_1$ to $\square_3$ control the overall distribution of the number of partnerships, and the parameter $\square$ determines the relative prevalence of the types of partnerships. The parameters $\square_1$ and $\square_2$ control the length of the two types of partnerships. Together, these parameters control the amount of concurrency and the lengths of the overlaps.

We simulate one partnership history at a time, and this history can be conceptually represented as the joint history of three possible partners at each point in time:

	Partner 1			Partner 2			Partner 3		
DAY	**Status**	**Change**	**Type**	**Status**	**Change**	**Type**	**Status**	**Change**	**Type**
1	1	0	1	0	1	2	0	0	0
2	1	0	1	1	0	2	0	0	0
3	1	0	1	1	0	2	0	0	0
4	1	0	1	1	-1	0	0	0	0
5	1	0	1	0	0	0	0	0	0

Each row represents a day, *status* = 1 if a partnership is present and 0 otherwise, *change* = 1 when a partnership is added, -1 when a partnership is ended and 0 otherwise, and *type* = 1 for a spousal partnership, 2 for a non-spousal partnership and 0 otherwise. The value of *change* is assessed each day using a random number generator:

Losing a partner:

1. Generate a random number from a uniform(0,1) distribution, call it RUNIF.
2. If {type = 1 and RUNIF < $\square_1$} or {type =2 and RUNIF < $\square_2$} end partnership
   else partnership continues.

Gaining a partner:

1. Determine the current number of partners, call it *num.*
2. Generate a random number from a uniform(0,1) distribution, call it RUNIF.
3. If {partners=*num* and RUNIF < $\square_{num+1}$} add an additional partner.

And the type of partner gained is assigned in a similar fashion:

1. Generate a random number from a uniform(0,1) distribution, call it RUNIF
2. If {RUNIF < $\square$} the type is spousal
   else the type is non-spousal.

So for the *ith* partnership on the *jth* day, the change in state, ☿$_{ij}$, is given by:

$$\square_{ij} = \begin{cases} 1 & \text{if RUNIF } \square\, \square_{num+1} \\ \square 1 & \text{if (RUNIF } \square_1 \text{ and } T_{ij} = 1) \text{ or (RUNIF } \square_2 \text{ and } T_{ij} = 2) \\ 0 & \text{otherwise} \end{cases}$$

and the type of partnership *i* on the *jth* day, T_{ij} is given by:

$$❄_{ij} = \begin{cases} ❄_{ij-1} & \text{if } \square_{ij} \square 0 \\ 1 & \text{if } (\square_{ij} \square 1 \text{ and RUNIF } \square \square) \\ 2 & \text{if } (\square_{ij} \square 1 \text{ and RUNIF } \square \square) \\ 0 & \text{if } \square_{ij} \square \square 1 \end{cases}$$

Using this procedure, we generate 200 years of sexual history (by day) for each of 1000 cases. To assign the observation interval covered by the survey, we need to establish an interview date and the sexual age (number of years sexually active) for each case. To ensure that the simulation process has reached equilibrium, we begin at year 50, and move up to the next observed start of a partnership. This is defined as the start of the observation interval. We then sample with replacement from the distribution of sexual age (percentiles are shown in Table A1b in the appendix), and move forward that number of years. This is defined as the end of the observation interval. This process generates a distribution of cases with the same sexual age distribution as the observed data. Each sexual history begins with the start of a sexual relationship, but as the simulated interview date is independent of partnership status, the simulated history may end in the middle of one or more partnerships, censoring their end dates. Finally for each case, we retain information on at most the last three relationships. The resulting data set of 1000 cases we call the "true" data.

The values of the basic partner formation, dissolution, and type allocation parameters were chosen so that the partnership outcomes matched, as closely as possible, the observed estimates from the Ugandan data. These input values are shown in Table A1a of the Appendix.

Adding Error

We then add two types of error to the "true" data: reporting unit heaping, and recall error.

Reporting Unit Heaping

In the actual survey, respondents chose their own reporting units for each date. For example, if she reported that partnership 1 started 20 years ago and ended 3 weeks ago, the values and the units were recorded by the interviewer. This led to a mix of reporting units – days, weeks, months and years – and one form of reporting error. It is clear that an answer of "2 years" is a rounded version of the true date, and the error induced by this is the rounding is what we seek to capture. From the truncation of the lesser unit distributions (days, weeks and months), we can infer a rounding process. For example, consider the distribution of days observed for respondents who reported in days, summarized in the table below.

Reported Days	Frequency	Cumulative Percent
0	123	11.7
1	343	44.2
2	227	65.7
3	133	78.3
4	89	86.7
5	47	91.2
6	31	94.1
7	8	94.9
8	20	96.8
9	10	97.7
10	9	98.6
11-33	15	100.0

If start and end days are more or less randomly distributed with respect to interview date, then we would expect a fairly uniform distribution of day frequencies. Instead, we find a severely right-skewed distribution. By 3 days, over 2 thirds of those who are going to report in days have done so, and by 7 days, respondents have clearly made the reporting transition to weeks. Similar patterns can be observed in the other reporting units.

Based on these patterns we developed a mixture distribution for the reporting units across the time line. It is loosely based on the observed data. Parts of the scale are reported exclusively in days, weeks, months or years, but an exponential decay in the probability of the smaller reporting unit then guides the transition to the next larger unit. The effect on the days distribution, for example, is that days in the interval from 0 to 3 are assigned exclusively to the day reporting unit (and are thus unchanged), days in the interval between 4 and 10 are increasingly defined in weeks, and after 10 days reporting is exclusively in weeks (until 14 days, where the month mixture begins). Notice that 4 days is the 78.3 percentile and 10 days is the 98.6 percentile. We use similar percentile cutoffs for the other reporting units.

The mixture distributions were determined using an exponential survival function. Take as an example the first mixture for the interval of 4 to 10 days. The survival function $S(t)$ has the constraints of $S(3)=1$ and $S(11)=0$. Thus when a date is 3 days, the unit choice is always days and when a date is 11 days, the unit choice is always weeks. The function *S(t)* is defined in (1), where the value of [illegible] is determined by the range of the mixture interval.

$$S(t)\ \square\ e^{\square\square(t\,\square)} \qquad (1)$$

A similar approach was taken for the other mixtures. A graphical representation of these functions is shown in Figure 2, and the cutpoints for the mixture interval are shown in Table A3 in the appendix.

The following procedure was used to assign the units for each date. The "true" date was compared to the first column in Table A3. If the date fell outside of the intervals listed there (into an interval not requiring a mixture), it was converted to the appropriate units, rounding to the nearest integer. Otherwise, we calculate *S*(date) from (1) for the appropriate mixture distribution, and generate a Bernoulli trial with probability of success equal to *S*(date). If the trial is a success, we express the date in the finer units, otherwise the coarser.

Recall Error

To account for recall mistakes, additional errors were added to the exact starting and ending dates for all dates greater than 365 days. These errors were assumed to be normally distributed, with a mean of 0 and a variance dependent on the date. We assume that the farther back in time the date is, the greater the spread in the error distribution. For the exact time of 12 months, the variance was determined by setting the 5^{th} percentile equal to 6 months and the 95^{th} percentile equal to 18 months. This results in a standard deviation of .3040. For the exact time of 10 years, the 5^{th} and 95^{th} percentiles were set equal to 8 years and 12 years, respectively. This results in a standard deviation of 1.2158. Linear interpolation was used to assign the variance for the intervening years[2], and the 10-year value was used for all subsequent years.

For spousal partnerships, the errors associated with the start and end dates were sampled independently. In a handful of instances, the addition of errors to the starting and ending dates resulted in the starting date less than the ending date. When this happened, the errors were resampled until the ending date was less than the starting date. For non-spousal partnerships, this

[2] Specifically, the equation is $\square^{\square}$ = 0.1013*(days/365) + .2028.

approach had to be modified[3]. Here we assigned the error associated with the end date as described, but then took the "true" difference, added it back to the end date, and rounded the result to the nearest appropriate reporting unit.

FINDINGS

The key question is whether the error in the date distributions changes our estimates of the prevalence and attributes of concurrency. We will look at three measures: (1) the prevalence of concurrency; (2) the length of overlap given concurrency; and (3) the types of partnerships that are found to be concurrent.

The prevalence of concurrency is measured in four ways here: concurrency between partners 1 and 2 (that is, most recent and second most recent); partners 1 and 3; partners 2 and 3; and finally any concurrency at all in the last 3 partners. The estimates, based on the simulated "true" dates and on the dates with added error are shown in Table 1. For all of these measures, the measurement error increases our estimates of concurrency. The error is smallest for the most recent two partners, and largest for partners farthest away in time. If there were no bias in the type of error added to the true dates, or if the bias were added equally to both start and end dates, then we might expect to find no systematic difference in the estimated concurrency measures. But clearly this is not the case. Overall, the effect is not large: if the true prevalence of at least one instance of concurrency in the last three partners is 33.3%, the measurement error raises this estimate to 38.6%: an increase of about five percentage points, or 16%. But the consistency of the overestimates indicates that the addition of error increases the rate of false positives (concurrency identified where it does not exist) more than the rate of false negatives.

This can be clearly seen in the classification tables shown in Table 2. Here we crosstabulate the "true" concurrency status of each case by the status after error has been added. The cells represent the number of cases in each category. The proportion of misidentified cases ranges from 9% to 15%, with the highest level again among partnerships farthest away in time. Overall, the proportion misidentified is 10%. The rate of false positives is higher than false negatives in the first and last concurrency combination. While the false negative rate is higher for

[3] These partnerships were typically less than 2 years long, and it was not uncommon for the independently generated errors to reverse the order of the start and end dates. Here, resampling would have led to biased errors.

the middle combination (of partner 1 and 3), the number of cases here is quite small, so its overall contribution to the positive bias in concurrency identification is minimal.

To understand the reason for this upward bias caused by the measurement error, it is important to rule out any simulation artifacts – in particular, to determine whether the error we have added is unbiased. With respect to the recall error, no bias is added to the true date, as the errors associated with each year are normally distributed with mean 0. With respect to the reporting unit heaping error, however, a subtle bias is introduced. Unit heaping has a general bias of increasing the value of the date (where increase refers to placing it farther back in time). Four days becomes a week, 3 weeks becomes a month, and 8 months becomes a year. However, to establish concurrency, we compare the starting date of the most recent partnership with the ending date of the prior partnership. If these two are close enough to potentially be misclassified, then they are often expressed in similar time units. As a result the *relevant* error is essentially unbiased[4], and the upward bias we observe in the concurrency estimate must a function of something else.

That "something else" is most likely the distribution of waiting times between partners that do not overlap, and crossover in start and end dates[5] between partners that do. The average median number of months spent in each state is shown in Table 3. False positives only occur when partnerships do not truly overlap. In this case, we can see that the waiting time between partnerships tends to be fairly short (at least for the first and last partnership combinations): on the order of 8 months. This is driven by the basic parameters of partner formation, and the values chosen here clearly ensure that persons are almost always involved in a partnership. For a false positive to occur, then, it is only necessary that the cumulative error be on the order of 8 months, not very large. On the other hand, for false negatives to occur, a true crossover between two partnerships must be eliminated. In this simulation, the median monthly crossover is fairly long: on the order of 40 months. To eliminate this crossover, the error has to equally large. Thus, for a uniform error distribution, the probability of an error large enough to eliminate a waiting interval

[4] Note that this bias does typically matter for estimates of relationship length. If the bias were added equally to the start and end dates of all relationships, there would be no problem. But end dates are more recent than start dates, so the units they are reported in are smaller, and more accurate. The start dates, by contrast, are more upwardly biased. The net effect is to bias the length estimates upwards.

[5] Note that crossover is slightly different than overlap. Overlap refers to the time that the two partnerships are held concurrently, while crossover refers to the difference between the start of the current partnership and the end of the

is much more likely than one large enough to eliminate a crossover. The result is that we get a higher rate of false positives – except for the partner 1 and 3 combination, where the median waiting interval is unusually long. We will take up the generalizability of this finding in the conclusion.

The next attribute of interest is the estimated overlap of the two partners, given that concurrency is identified. The mean (and median) overlaps in months for the three possible concurrency combinations is given in Table 2. The measurement error consistently exerts a downward bias on the estimate of overlap, with the greatest bias occurring for the partners farthest back in time. Given the increased levels of concurrency found above, this makes sense. In general, the additional "misidentified" concurrent partners will have a short overlap appear in their start and end dates with the addition of the measurement error. As overlaps are otherwise fairly long among the true concurrent partners (from 2 to 4 years on average), the effect of adding these new misidentified cases will be to lower the mean and median overlap.

Finally, we are interested in the types of partnerships that are identified as concurrent. There are four different pairing combinations in this simulation: spouse-spouse, spouse-nonspouse, nonspouse-spouse, and nonspouse-nonspouse. The order matters, because it represents which partnership was added more recently. The distribution of these different pairings for the "true" and error-added datasets can be seen in Table 4. As we have seen in with the other measures, the largest errors occur with partnerships that are further back in time. Net of this, there does not appear to be any systematic pattern in which types of pairings are erroneously misclassified, as either false positives or false negatives.

Note that in the partner 1 and 2 concurrency, it is very common for the two partners both to be spouses, and in all of the concurrencies, the majority of pairings find at least one spouse . This plays an important role in determining the length of the crossovers that we saw above, as spousal relationships are fairly long term. This was not a "target" value used in setting up the simulation, and, in fact, this is not the pattern we observe in the Ugandan data, where only about 25% of the pairings are spouse-spouse in the partner 1 and 2 concurrency, and less in the other cases. We did not anticipate the role that this would play in our estimation, and will therefore need to reconsider the simulation rules if we wish to represent the observed data more accurately.

---

prior partnership. When the prior partnership is wholly contained in the current partnership interval, then the crossover measure will typically be larger than the overlap.

DISCUSSION

Measurement error appears to increase survey estimates of concurrency for the behavioral pattern examined here, though the amount of error is relatively small. The upward bias generated by the measurement error is largely due to a specific behavioral pattern: the small intervals that persons spend in the single state, compared to the large intervals in the crossovers if the partnerships overlap. The length of the crossovers in this simulation is driven by the high rates of concurrency involving two spouses, especially when partnerships 1 and 2 are concurrent. As this pattern does not accurately represent our observed data, one should be careful in drawing implications from the rates of error and misclassification drawn in this simulation study. In particular, there is a good chance that we have overestimated the impact of measurement error on the estimates of concurrency. What is clear is that waiting intervals and crossovers are important parameters in this process, and attempts to estimate the measurement error effect will have to be conditional on the values observed for these parameters in different contexts.

Net of this, we can say that the magnitude of the bias introduced depends on what is being estimated. Simple aggregate measures of concurrency prevalence and length of overlap will fare pretty well. We observed a small upward bias of about five percentage points on a true prevalence of 33 percent, and a small downward bias typically under 10% for the mean overlap. On a casewise basis, the misclassification rate (which combines false negatives and false positives) ranged from 9-15%, with an overall rate of 10%. The pairings of concurrent partners are differentially affected, with no systematic bias apparent in our simulation.

In general, the farther back in time that these partnerships occur, the larger the measurement error effects. This is clearly due to the way in which we have modeled the variance in the recall error, as we have assumed that it increases linearly over time. Other models for the variance in the recall error can easily be incorporated into this modeling exercise, but the effects are predictable.

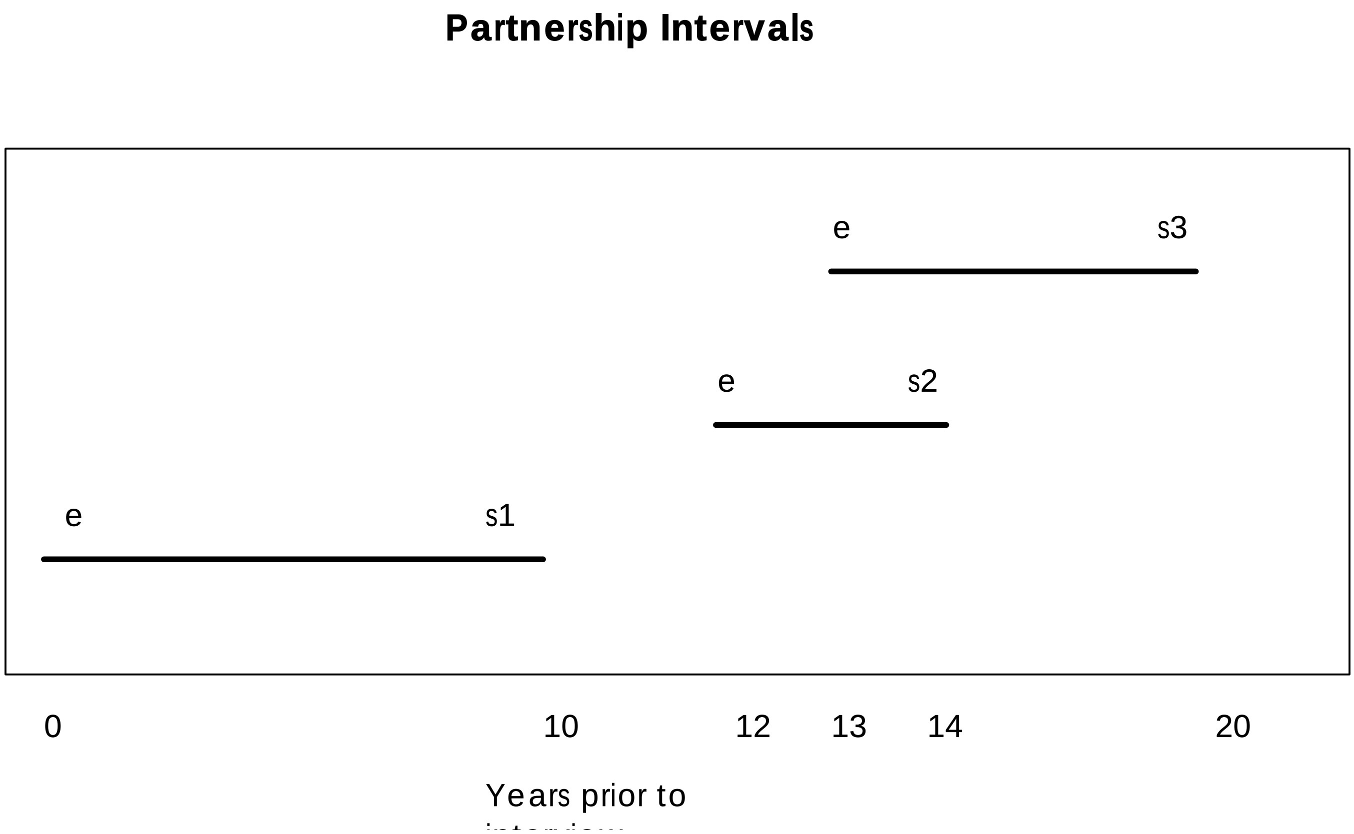

Figure 1. Graphic representation of the dates of first and last sex that a respondent might report for each of their three last partners. The x-axis represents the time line with 0 as the date of the interview, and positive numbers showing increasingly farther back in time. The end dates of the relationships are shown by "e1" , "e2" and "e3", for the most recent, second most recent, and third most recent partners respectively, and the corresponding start dates are shown by "s1", "s2", and "s3". In this example, partners 2 and 3 are concurrent for about

Table 1. Levels of Concurrency

Concurrency with:	**Simulated "true"**	**Simulated w/error**	**Percentage Point Difference**
Partners 1 and 2	30.5	34.0	+3.5
Partners 1 and 3	4.1	6.4	+2.3
Partners 2 and 3	9.7	14.1	+4.4
Any of the above	33.3	38.6	+5.3

Table 2. Classification Table

Simulated "true"	**Simulated w/error**		**% Misclassified (false pos / false neg)**
Partners 1 and 2	*not concurrent*	*concurrent*	
not concurrent	271	45	8.9
concurrent	10	295	(14.2 / 3.3)
Partners 1 and 3	*not concurrent*	*concurrent*	
not concurrent	319	31	9.7
concurrent	7	34	(8.9 / 17.1)
Partners 2 and 3	*not concurrent*	*concurrent*	
not concurrent	242	52	15.3
concurrent	8	89	(17.7 / 8.2)
Any 2 partners	*not concurrent*	*concurrent*	
not concurrent	229	59	10.4
concurrent	6	327	(19.6 / 1.8)

Table 3. Median months of waiting intervals and crossovers

Between:	**Waiting Interval**	**Crossover†**
Partners 1 and 2	8	39
Partners 1 and 3	26	41
Partners 2 and 3	7	42

† Crossover refers to the difference between the start of partnership1 and the end of partnership 2 in cases where the start of partnership 1 is farther back in time.

Table 4. Months of overlap given concurrency: Mean (Median)**†**

Overlap with:	**Simulated "true"**	**Simulated w/error**	**Percent Difference**
Partners 1 and 2	48 (24)	46 (24)	-4.2 (0)
Partners 1 and 3	22 (14)	20 (12)	-9.1 (16.7)
Partners 2 and 3	25 (17)	23 (12)	-8.0 (29.4)

†Overlap refers to the period that the two partnerships are held concurrently, and will be lower than the crossover period due to short partnerships that lie wholly within another partnership interval.

Table 5. Pairing Combinations of Concurrent Partners

	Pairing combination (1:2)†			
	s:s	**s:n**	**n:s**	**n:n**
Partners 1 and 2				
"true"	43.0	34.8	14.1	8.2
error	39.7	37.6	12.9	9.7
% pt difference	3.3	-2.8	1.2	-1.5
Partners 1 and 3				
"true"	17.1	75.6	0.0	7.3
error	20.0	61.5	3.1	15.4
% pt difference	+2.9	14.1	+3.1	-8.1
Partners 2 and 3				
"true"	19.6	58.8	5.2	16.5
error	17.0	47.5	7.1	28.4
% pt difference	-1.4	-11.3	+1.9	+11.9

† The column labels X:Y here refer to the *more recent:less recent partners*, where *s* refers to a spouse and *n* refers to a non-spouse.

APPENDIX: SIMULATION DETAILS

Table A1. Input parameters.

1a. Partner formation, dissolution and type

Parameter	Description	Value in Simulation
$\square_1$	P(gain one partner | currently 0 partners)	$4.9451 * 10^{-3}$
$\square_2$	P(gain one partner | currently 1 partner)	$3.9246 * 10^{-4}$
$\square_3$	P(gain one partner | currently 2 partners)	$7.0644 * 10^{-5}$
$\square_1$	P(lose a spousal partner in a given interval)	$1.1682 * 10^{-4}$
$\square_2$	P(lose a non-spousal partner in a given interval)	$1.1774 * 10^{-3}$
$\square$	P(a partner is spousal)	0.4249

1b: Years of Sexual Activity

Years Sexually Active	Percentile
6	25.6
11	50.5
19	74.9
26	90.1
29	94.5
33	98.9

Table A2. Target output values from observed data

	Type		Mean Length	
	Observed	Simulated	Observed	Simulated
Partner 1				
Spousal	70.4	65.9	100.6	100.8
Non-spousal	29.6	34.1	27.3	26.5
NA	-	-	-	-
Partner 2				
Spousal	29.1	24.9	89.1	129.6
Non-spousal	50.5	37.2	21.7	23.6
NA	28.8	37.9	-	-
Partner 3				
Spousal	9.5	8.5	85.8	74.2
Non-spousal	35.9	30.6	22.4	27.5
NA	54.6	60.9	-	-

Table A3. Mixture intervals for assigning reporting unit heaping

"True" Day Interval	Mixture	Finer Unit Cut Points	Observed Percentile
4-10	Days and Weeks	4 days	78.3
		10 days	98.6
15-31	Weeks and Months	2 weeks	76.0
		4 weeks	97.0
181-365	Months and Years	6 months	84.4
		12 months	94.7
510-570	18 Months and Years	†	

† There was significant additional heaping observed in the data on the point 18 months. To preserve this feature, we use a mixture of 18 months and 1 or 2 years for dates in the interval from 510 to 570 days.

REFERENCES

Kretzschmar, M, and M Morris. 1996. "Measures of concurrency in networks and the spread of infectious disease." *Math. Biosc.* 133:165-95.

Morris, M, and M Kretzschmar. 1997. "Concurrent partnerships and the spread of HIV." *AIDS* 11:641-8.

Wawer, M J. 1993. "Ugandan Sexual Network/Behaviors Study for HIV Prevention." *National Institutes of Health Research Grant*.