Practice of Epidemiology

Dynamic Variation in Sexual Contact Rates in a Cohort of HIV-Negative Gay Men

E. O. Romero-Severson*, E. Volz, J. S. Koopman, T. Leitner, and E. L. Ionides

* Correspondence to Dr. E. O. Romero-Severson, Theoretical Biology and Biophysics Group, Los Alamos National Laboratory, P.O. Box 1663, Mailstop K710, Los Alamos, NM 87545 (e-mail: eoromero@lanl.gov).

Initially submitted July 28, 2014; accepted for publication February 9, 2015.

Human immunodeficiency virus (HIV) transmission models that include variability in sexual behavior over time have shown increased incidence, prevalence, and acute-state transmission rates for a given population risk profile. This raises the question of whether dynamic variation in individual sexual behavior is a real phenomenon that can be observed and measured. To study this dynamic variation, we developed a model incorporating heterogeneity in both between-person and within-person sexual contact patterns. Using novel methodology that we call iterated filtering for longitudinal data, we fitted this model by maximum likelihood to longitudinal survey data from the Centers for Disease Control and Prevention's Collaborative HIV Seroincidence Study (1992–1995). We found evidence for individual heterogeneity in sexual behavior over time. We simulated an epidemic process and found that inclusion of empirically measured levels of dynamic variation in individual-level sexual behavior brought the theoretical predictions of HIV incidence into closer alignment with reality given the measured per-act probabilities of transmission. The methods developed here provide a framework for quantifying variation in sexual behaviors that helps in understanding the HIV epidemic among gay men.

disease transmission; gay men; HIV; HIV risk; iterated filtering; partially observed Markov process; sexual behavior

Abbreviations: HIV, human immunodeficiency virus; LOESS, local regression scatterplot smoothing.

Although the number of new human immunodeficiency virus (HIV) infections in the United States remained approximately stable from 1991 to 2006, the incidence rate in young gay men increased by 58% over the same period (1). The increase in transmission rates is difficult to explain in the context of a virus with a low per-act probability of transmission (2–4) and the availability of effective modes of prevention, such as condoms and antiretroviral medication (5).

Vittinghoff et al. (6) found that unprotected receptive anal sex results in HIV infection at a probability of 0.0028 per contact. Standard epidemiologic theory based on homogeneous, well-mixed populations (7) predicts that an HIV prevalence of 21%, which is consistent with estimated prevalence among men who have sex with men (8), could be reached only if each individual had 15 new sexual partners per year for an average of 30 years. Such a high level of homogeneous, sustained high-risk behavior seems improbable.

The inconsistency between rising rates of man-to-man HIV transmission and observed contact rates and transmission

probabilities has been addressed by decades of theoretical work, including analysis of risk groups (9–11), enduring partnerships (12–14), natural history of infection (15, 16), and variability in the level of risky sexual behavior over time, termed episodic risk (10, 17–19). All of the above factors, excluding episodic risk, have been empirically well characterized. Resolving whether or not episodic risk is a real, measurable phenomenon will help epidemiologists to parameterize more realistic models of HIV transmission dynamics.

A conceptually simple typology of HIV risk factors that illustrates the issues at play divides along 2 basic axes: population versus individual and between-units versus within-units. For example, the prevalence of intravenous drug use (an aggregate feature) in Sweden versus Latvia (prevalence between 2 populations) helps to explain the differential spread of HIV in those populations (20, 21). Between-person risk factors for the spread of HIV are well known: anal sex, unprotected sex, and birth to an infected mother, for example. Considering all possible partners, a person who never has unprotected receptive

anal intercourse is at lower risk than someone who does. However, the behaviors of individuals and whole populations are not constant over time, producing within-unit variation over time.

In this paper, we aim to illustrate a broadly applicable statistical framework for estimating both within-person and between-person episodic risk parameters from a longitudinal data set of male-male sexual contacts. Our method relates an underlying model of dynamic variation in sexual contact rates to aggregate count data in variable observation periods. A methodological innovation we illustrate in this paper is the application of iterated filtering methods (22, 23), originally developed for investigating mechanistic models using time series data, to fit models in a longitudinal study with data consisting of many short time series. Finally, we demonstrate the influence of episodic risk on epidemic dynamics by simulating potential epidemic trajectories from a simple stochastic model.

METHODS

Sexual contact data

The data in this study came from a large cohort of HIV-negative gay men in 3 US cities (San Francisco, California; Denver, Colorado; and Chicago, Illinois) who reported having had either anal sex or receptive oral sex with ejaculation in the previous 6 months. The data were collected beginning in 1992, as described by Vittinghoff et al. (6). The men were asked about the numbers and types of their sexual contacts over the course of 3 follow-up interviews. A total of 882 men completed the initial interview and 3 follow-up interviews, for a total of 2 years of observation per person divided into four 6-month observation periods.

We aggregated the count data into broadly relevant behavioral categories: total number of sexual contacts (any type), anal acts, oral acts, protected sex, unprotected sex, sex with a partner of unknown HIV status, and sex with a partner believed to be HIV-positive. For example, the anal sex category included all reported instances of anal sex, regardless of whether or not they were protected, insertive, or with an HIV-positive partner. However, the only type of oral contact measured in the study was unprotected receptive oral sex with ejaculation.

The data did not differentiate between contacts made with long-term partners and those made with short-term partners, such that we could not disentangle the rate of partnership formation from the rate of sexual contact itself.

A descriptive model of contact rates that accounts for between-person and within-person variation

Romero-Severson et al. (24) identified 4 features of contact rate data that should be included in a statistical model of contact rates: 1) heterogeneity between individuals, 2) heterogeneity within individuals over time, 3) individual-level autocorrelation, and 4) secular trend over the duration of the cohort study. Figure 1 shows plots of these 4 factors in the data. To model these phenomena, we suppose that each individual has a latent rate $X_i(t)$ of making contacts of a specific type. Each data point, y_{ij} , is the number of reported contacts for individual i between time t_{j-1} and time t_j , where $i = 1, \dots, 882$ and $j = 1, \dots, 4$. The unobserved process $\{X_i(t)\}$ is connected to

the data through the expected number of contacts for individual i in reporting interval j , which we write as

$$C_{ij} = \alpha^{j-1} \int_{t_{j-1}}^{t_j} X_i(t) dt, \quad (1)$$

where α is an additional secular trend that accounts for the observed decline in reported contacts. A basic stochastic model for homogeneous count data would model y_{ij} as a Poisson random variable with mean and variance equal to C_{ij} (25). However, the variance in the data is much higher than the mean value of the data distribution (24). To account for this, we assume that the data are negative binomially distributed (26, 27), which is a generalization of a Poisson distribution that allows for increased variance for a fixed mean. This leads to the model

$$y_{ij} \sim \text{NegBin}(C_{ij}, D_i), \quad (2)$$

with mean C_{ij} and variance $C_{ij} + C_{ij}^2/D_i$ (probability mass function defined in Web Appendix 1, available at <http://aje.oxfordjournals.org/>). Here, D_i is called the dispersion parameter, with the Poisson model being recovered in the limit as D_i becomes large. The dispersion, D_i , can model increased variance in comparison with the Poisson distribution for individual contacts, but it does not result in autocorrelation between measurements on an individual over time, which is observed in the data. To model this autocorrelation, we suppose that individual i has behavioral episodes within which $X_i(t)$ is constant but the individual enters new behavioral episodes at rate R_i . At the start of each episode, $X_i(t)$ takes a new value drawn from a gamma distribution with mean μ_X and variance σ_X :

$$X_i(t) \sim \text{Gamma}(\mu_X, \sigma_X). \quad (3)$$

To complete the model, we also assume gamma distributions for D_i and R_i :

$$D_i \sim \text{Gamma}(\mu_D, \sigma_D), \quad (4)$$

$$R_i \sim \text{Gamma}(\mu_R, \sigma_R). \quad (5)$$

The parameters σ_X , σ_D , and σ_R control individual-level differences in behavioral parameters, allowing the model to encompass a wide range of sexual contact patterns.

Figure 2 shows one possible realization of the model for a single individual. The distinction between the effects of the rate at which new behavioral episodes begin, R_i , and the dispersion parameter, D_i , is subtle, since both model within-person variability. The signal in the data about distinct behavioral episodes could be overwhelmed by a high variance in the number of reported contacts resulting from a low value of D_i . Whether the data are sufficient to identify both R_i and D_i is an empirical question that we address below.

Likelihood-based parameter estimation and model selection

Our stochastic dynamic model for contacts was a partially observed Markov process (Web Appendix 2; see also Bretó et al. (28)). Likelihood-based inference was carried out using

Figure 1. Features of a set of longitudinal data on rates of sexual contact among human immunodeficiency virus (HIV)-negative gay men in the United States, Centers for Disease Control and Prevention Collaborative HIV Seroincidence Study (1992–1995). A) Secular trend from the time of enrollment in the cohort; B) average rate of sexual contact per month; C) rates of sexual contact over time; D) bias-corrected autocorrelation (Web Appendix 1). Black bars show autocorrelation >0 , while gray bars show autocorrelation ≤ 0 . The mean (0.076) and standard error (0.0094) of the autocorrelation imply a small but positive autocorrelation (95% confidence interval: 0.057, 0.094).

iterated filtering (22) implemented in *pomp*, version 0.43-4 (29), running in R2.15.3 (30). Iterated filtering is a Monte Carlo algorithm which computes the maximum likelihood estimate for partially observed Markov process models. Filtering is the numerical computation of estimating unobserved states and evaluating the likelihood function for a partially observed Markov process. Iterated filtering carries out multiple filtering operations using a sequential Monte Carlo filter, with perturbations in the unknown model parameters designed so that successive filtering operations converge

toward the maximum likelihood estimate. The sequential Monte Carlo method is a flexible nonlinear non-Gaussian filtering method, also known as the *particle filter* (31), in which the unknown distribution of the latent dynamic variables is represented by a Monte Carlo sample from this distribution (known as a *swarm of particles*). Successive iterations of the filtering process make successively smaller perturbations to the parameters, with the heuristic that the optimization process is cooling toward a freezing point which is theoretically guaranteed to be a local maximum of the likelihood function

Figure 2. A model of rates of sexual contact that accounts for between-person and within-person variation in sexual behavior over time, Centers for Disease Control and Prevention Collaborative HIV Seroincidence Study (1992–1995). The graph shows one possible realization of the model for a single individual, i . The length of behavioral intervals (horizontal lines) over which an individual's contact rate (left axis) is constant is controlled by R_i , having the mean value $1/R_i$. At the beginning of a new behavioral interval, a new contact rate is drawn from a gamma distribution mean, μ_X , and standard deviation, σ_X . The box plots illustrate the distribution of the possible number of sexual contacts (right axis) for the given pattern of behavioral intervals shown over each 6-month observation period. The positive latent variable D_i governs the dispersion of the actual number of contacts within a given behavioral interval. As D_i decreases toward zero, the number of contacts observed in each observation period will be more variable. HIV, human immunodeficiency virus.

(23). Increasing the number of particles and increasing the duration of the cooling process may improve the practical performance of the algorithm, while adding to the computational expense.

All optimizations were calculated by running sequential, fast-cooling particle filters (200 particles, 100 iterations, cooling factor 0.95). At the end of each optimization, the log-likelihood was estimated by a running a single-pass particle filter with 10^5 particles. A new particle filter was initialized at the current best parameter set and run using the same fast-cooling strategy.

Profile likelihood methods were used to compute maximum likelihood estimates and confidence intervals for each type of sexual contact by means of the selected model (Table 1, Web Figures 1–9). The process of applying sequential particle filters was stopped once the profiles appeared to be unchanged by further applications of the particle filter. The maximum likelihood estimates reported in Table 1 were found by fitting a local regression scatterplot smoothing (LOESS) curve to the profile likelihood and finding its maximum (32). The confidence intervals were obtained by finding the points on either side of the LOESS curve that were 1.92 log-likelihood units below the maximum likelihood estimate. The reported log-

likelihood was the maximum of the likelihoods for each profile for a given contact type. Nested models can be compared using likelihood ratio tests (33) and nonnested models by Akaike's Information Criterion (34). Either way, evidence for a parameter is strong if addition of that parameter increases the maximized likelihood by much more than 1 log unit.

Simulation from a simple epidemic model with time-variable contact rates

Simulations of the epidemic dynamics assuming the time-variable contact rates implied by our analysis were coded as Markov chains in R3.0.3 (30). Individuals were assumed to be either susceptible (S) or infected (I) and to select partners randomly proportional to their contact rates. At entry into the sexually active population, individuals were assigned a contact rate $X_i \sim \text{Gamma}(\mu_X, \sigma_X)$ and an overdispersion parameter $D_i \sim \text{Gamma}(\mu_D, \sigma_D)$. To simulate the trajectory of the system, we evaluated the probability of each event over a period of 1 month. Given that the model parameterization implies much less than 1 infectious contact per month, this time step is reasonable. At each step we determined, in order: 1) which individuals became infected; 2) which individuals started a new behavioral episode; 3) which individuals died or were otherwise removed; and 4) how many new individuals entered the system. In our simulations, each susceptible individual makes $Q_i \sim \text{NegBin}(X_i, D_i)$ contacts per month. Of those, $\text{Binomial}(Q_i, \eta_i \beta / (\eta_S + \eta_i))$ contacts are infective, where $\eta_S = \sum_{j \in S} Q_j$ is the total number of contacts made by susceptible individuals, $\eta_i = \sum_{j \in I} Q_j$ is the total number of contacts made by infected individuals, β is the per-act probability of transmission, and $\text{Binomial}(n, p)$ is the binomial distribution with mean np and variance $np(1-p)$. A susceptible individual who has 1 or more infective contacts in a month becomes infected. Each individual redraws a new $\text{Gamma}(\mu_X, \sigma_X)$ contact rate each month with probability $1 - \exp(-\mu_R)$. Susceptible and infected individuals are removed with probability $(1 - \exp\{-\zeta\})$ and $(1 - \exp\{-(\zeta + \delta)\})$ per unit of time, respectively, where ζ is the general rate of removal and δ is the death rate for infected individuals. In each unit of time, a constant number of new susceptible individuals is added (i.e., the smallest integer greater than κ/ζ) so that the long-run, infection-free population size is approximately κ . In all simulations, $\kappa = 3,000$, $\zeta = 1/(40 \times 12)$ contacts/month, and $\delta = 1/(10 \times 12)$ contacts/month.

RESULTS

Model selection

Table 2 shows strong evidence (>50 log units of improvement in the likelihood) for each parameter excluding σ_R (between-person heterogeneity in the length of a behavioral episode). The full model with σ_R is no better than the model without σ_R . With perfect likelihood maximization, model 6 can only have a higher maximized likelihood than model 5, since it contains model 5 when $\sigma_R = 0$. We explain the lower maximized likelihood for model 6 to be a result of the numerical difficulty of maximization for this more complex model, combined with no substantial statistical benefit from including

Table 1. Maximum Likelihood Parameter Estimates and Corresponding Likelihoods for 9 Types of Sexual Contact, Centers for Disease Control and Prevention Collaborative HIV Seroincidence Study, 1992–1995

Type of Contact	Model Parameter ^{a,b}												Log-Likelihood
	μ_X , months	95% CI	σ_X , months	95% CI	μ_D	95% CI	σ_D	95% CI	μ_R , months	95% CI	α	95% CI	
Total (all contacts)	1.75	1.62, 2.01	2.67	2.35, 3.20	3.81	2.76, 4.59	4.42	3.30, 5.03	0.04	0.03, 0.05	0.90	0.87, 0.93	–9,552.1
Anal	1.49	1.36, 1.64	2.33	2.05, 2.64	3.47	2.88, 4.01	4.15	3.38, 5.94	0.05	0.04, 0.06	0.90	0.87, 0.95	–9,092.2
Oral	0.23	0.20, 0.27	0.77	0.65, 1.03	1.07	0.84, 1.32	1.05	0.65, 1.57	0.02	0.01, 0.03	0.84	0.76, 0.90	–2,622.7
Insertive	0.93	0.85, 1.01	1.64	1.50, 1.94	2.26	2.00, 2.56	2.42	1.82, 3.35	0.03	0.03, 0.04	0.90	0.86, 0.94	–7,438.0
Receptive	0.51	0.44, 0.56	0.98	0.89, 1.17	2.14	1.80, 2.61	2.59	2.12, 3.44	0.03	0.02, 0.04	0.95	0.90, 1.0	–5,688.4
Protected	1.07	0.98, 1.18	1.91	1.73, 2.18	2.64	2.23, 3.11	2.90	1.85, 3.65	0.04	0.03, 0.05	0.95	0.89, 0.98	–8,179.9
Unprotected	0.40	0.33, 0.46	0.91	0.77, 1.20	1.19	0.95, 1.59	1.69	1.14, 2.49	0.04	0.03, 0.05	0.81	0.76, 0.89	–3,775.2
Positive ^c	1.62	1.01, 2.32	7.45	5.21, 9.78	0.99	0.68, 1.98	3.27	1.92, 5.08	0.09	0.06, 0.11	0.71	0.64, 0.84	–2,190.6
Unknown	1.08	0.98, 1.20	1.67	1.51, 1.90	3.19	2.74, 3.94	3.76	2.80, 5.86	0.04	0.03, 0.61	0.94	0.90, 0.98	–8,214.5

Abbreviations: CI, confidence interval; HIV, human immunodeficiency virus.

^a Parameters are defined in equations 1–5 (see text).

^b μ_D , σ_D , and α are dimensionless and have no units.

^c Sex with a partner believed to be HIV-positive.

this additional complexity. Therefore, for subsequent analysis, we adopted model 5 and fixed $\sigma_R = 0$.

Standard deviation of the number of contacts and between-person heterogeneity

The empirical sample mean and sample standard deviation were 1.53 contacts/month and 3.28 contacts/month, respectively. The maximum likelihood parameters are shown in Table 1 and Figure 3. All types of sexual contact showed high levels of between-person variability (high σ_X), suggesting long-term differences in individual contact rates. However, for every contact type, σ_X was lower than the empirical between-person standard deviation. For total contacts, σ_X was 19% lower than the empirical between-person standard deviation, and the confidence interval for σ_X did not include the empirical standard deviation. This discrepancy arises

because the raw standard deviation captures both within-person and between-person variance. Consider a population with no between-person variability in contact rates (in the long run, everyone has the same number of contacts) but a high level of within-person variance. If we observed that population for a limited period of time, we would see some persons with apparently high contact rates and some with low contact rates. The difference between those individuals is due solely to chance variance about a shared average contact rate. Studies that use the empirical standard deviation of contact rates to model between-person heterogeneity will overestimate its magnitude.

High levels of both between-person and within-person variability

The mean overdispersion (μ_D) is interpreted relative to Poisson variability. In the limit, as μ_D approaches infinity,

Table 2. Maximum Likelihood Parameter Estimates and Corresponding Likelihoods for 6 Models of Total Sexual Contacts (All Types), Centers for Disease Control and Prevention Collaborative HIV Seroincidence Study, 1992–1995

Model	Model Parameter ^{a,b}							Log-Likelihood
	μ_X , months	σ_X , months	μ_D	σ_D	μ_R , months	σ_R	α	
1	1.61		0.31				1	–10,288.3
2	1.73		0.6	0.67			0.99	–9,935.5
3	1.62	2.11	0.76				0.99	–9,772.9
4	1.71	1.81	1.43	1.71			0.99	–9,605.6
5	1.75	2.67	3.81	4.42	0.04		0.90	–9,552.1 ^c
6	1.79	2.63	3.98	4.91	0.04	0.01	0.93	–9,554.8

Abbreviation: HIV, human immunodeficiency virus.

^a Parameters are defined in equations 1–5 (see text).

^b μ_D , σ_D , and α are dimensionless and have no units.

^c The models show statistically significant improvement (increasing likelihood) though model 5.

Figure 3. Visual representation of the maximum likelihood parameter estimates for a model of variations in sexual behavior among human immunodeficiency virus (HIV)-negative gay men.

the number of observed contacts becomes Poisson-distributed with the rate determined by the integration of contact rates over behavioral intervals, C_{ij} . Decreasing μ_D increases the variance in the number of observed contacts without changing its mean. For each contact type (excluding HIV-positive and oral contacts), the average degree of overdispersion is relatively small (larger values of μ_D). However, the overdispersion parameter is highly heterogeneous in the population (high σ_D), implying the coexistence of persons with either highly stable (Poisson-like) or highly variable patterns of sexual contact over time.

For most contact types, the average duration of time for which an individual has a constant average contact rate is approximately 2 years ($\mu_R = 0.04$ contacts/month). Therefore, the statistical evidence for a rate $\mu_R > 0$ suggests that redrawing contact rates (i.e., episodic behavior) is allowing additional within-person variance that is not fully accounted for by μ_D and σ_D . Unlike the overdispersion described by μ_D and σ_D , this kind of behavioral change is on the appropriate time scale to correspond to structural life changes, such as formation of new partnerships.

Transmission models parameterized from cross-sectional versus longitudinal data

In the counterfactual case where the data set that we analyzed here was, in fact, a cross-sectional survey (i.e., collapsed over observational periods), we could still estimate between-person variability. In that case, we might consider the simple model in which individuals are characterized by a single contact rate drawn from a population distribution. To parameterize that distribution, we might consider the sample

Figure 4. Possible epidemic curves for the human immunodeficiency virus (HIV) epidemic under 3 parameterizations of variation in sexual behavior. The median values from 500 simulations are shown as lines, and the interquartile ranges (25th–75th percentiles) are shown as gray shading surrounding the 3 parameterizations. In the “homogeneous” case (dashed line), the epidemic was simulated with μ_X estimated by the sample mean (1.53 contacts/month) without any sources of between-person or within-person heterogeneity. In the “between heterogeneity” case (dotted line), the epidemic was simulated with μ_X estimated by the sample mean (1.53 contacts/month) and σ_X estimated by the sample standard deviation (3.28 contacts/month). In the “within + between heterogeneity” case (solid line), the epidemic was simulated with each parameter set to the estimated maximum likelihood estimate for total sexual contacts (all types). For all situations, the per-contact probability of transmission was set to 1/120, the average length of infection was set to 10 years, and the infection-free equilibrium population size was set to 3,000. The per-contact probability was selected such that the basic reproduction number in the “homogeneous” case was 1.53. In the “homogeneous,” “between heterogeneity,” and “within + between heterogeneity” cases, respectively, 239 of 500, 172 of 500, and 95 of 500 simulations died out before the 100-year mark.

mean (1.53 contacts/month) and sample standard deviation (3.28 contacts/month) as estimates of μ_X and σ_X , respectively. Figure 4 shows a comparison of the dynamics under 3 alternative parameterizations to illustrate that problems might arise when considering cross-sectional data alone. The dashed line illustrates random realizations of a stochastic model without any heterogeneity in which the contact rate is set to the sample mean monthly contact rate. The dotted line illustrates random realizations of a stochastic model in which the average contact rate is set to the sample mean monthly contact rate and the population standard deviation of contact rates is set to the sample standard deviation. The solid line shows the dynamics with the maximum likelihood parameters from the total contacts. In the context of this simple model, inclusion of episodic risk parameters estimated from the data greatly increases the prevalence of infection. In our simulations, after 35 years (omitting runs that died out), the average prevalence was 0.6% in the “homogeneous” case, 15% in the “between heterogeneity” case, and 20% in the

“within + between homogeneity” case (Figure 4), which is closer to the 21% reported for the US gay male population (8). Although this model is too simple to capture the full complexity of HIV transmission dynamics, it illustrates the potential role of heterogeneity within and between individuals in the theoretical explanation of observed HIV prevalence.

Contacts made with HIV-positive partners

All types of contacts except those made with a person known to be HIV-positive had an empirical mean close to μ_X . The empirical rate of contact with a person known to be HIV-positive (0.33 contacts/month) was much lower than the estimated value of $\mu_X = 1.62$ contacts/month. This discrepancy is due to the fact that contacts of this type showed a particularly high variability both within and between individuals. While many individuals reported consistent contact with an HIV-infected person over the course of the study, many also reported changes of more than 100 contacts between 6-month observation periods. This particular combination of both stable and highly variable contact patterns is difficult to model within our framework given the available data, making the results for this contact type less reliable.

DISCUSSION

In this paper, we have demonstrated application of a new general statistical framework for estimating both within-person and between-person variance in sexual behavior over time based on iterated particle filtering. This method provides a general approach for fitting arbitrary models to longitudinal behavioral data that could be expanded to include more complex and realistic behavioral models as more data become available. We showed that even in relatively short time series, we can discern the signature of significant within-person behavioral variability. Our model is general enough to account for both within-person and between-person variability for all observed contact types, excluding contacts made with a person believed or known to be HIV-positive.

There is a large body of work studying how behavior is modified in response to changing attitudes or perceived risk (35), such as serosorting of HIV-positive persons (36). That literature is largely focused on how otherwise static behaviors change in response to various external conditions. To our knowledge, our paper is the first to demonstrate empirically that sexual behavior can be highly variable in general—that is, that the assumption of static behavior is potentially problematic. The specific causes of the dynamic variation in sexual behavior we observed remain unknown.

Previous theoretical work has shown that episodic risk can greatly increase the rate of HIV transmission, especially from persons recently infected (10, 19, 18). A key effect of episodic risk observed in those papers was the maintenance of a high-risk susceptible pool that sustains the epidemic at higher levels than would otherwise be observed. In a population with only between-person heterogeneity, the highest-risk individuals (e.g., persons with high contact rates) become rapidly infected, while the remaining lower-risk individuals are more slowly infected. Episodic risk alters these dynamics by allowing

all individuals to experience brief episodes of high-risk behavior that leads to both elevated population risk and elevated acute-stage transmission. Episodic risk should be considered in the broader theoretical framework for understanding the rising HIV incidence among men who have sex with men in the United States.

Our analysis had important limitations that should be addressed in the collection and analysis of future data sets. First, the data were old (collection began in the early 1990s), which may limit their relevance to contemporary transmission dynamics. However, given the widespread availability of new treatment options in the 1990s that changed the perception of HIV risk, having a baseline estimate of how sexual behavior changes over time will provide context for future analyses. We did not find any newer studies in which investigators specifically reported counts of particular sex acts over time. Often, behavioral data were measured in terms of risk factors, such as whether the interviewee had had unprotected anal sex at least once in some previous time interval, rather than specific counts. Although we suggest caution regarding how these results are implemented in future models, we have demonstrated that episodic risk is a real and measurable phenomenon. Future studies will help clarify the contemporary nature of episodic risk.

The method developed in this paper could be readily expanded to other populations for which longitudinal behavioral data are available. The ideal data set would record the time of each contact and with whom the contact was made rather than collect aggregate counts over specified observation periods. Our method could be easily modified to obtain maximum likelihood estimates from potentially complex behavioral models with such data. Future work on episodic risk could include demographic and contextual covariates (such as age, relationship status, or drug use) that could be used to compile a more complete picture of how sexual behavior changes over the course of a person's lifetime. We recommend that future modeling studies for pathogens with long infectious periods include the potential for within-person variability in behavioral parameters, and we emphasize that, despite its cost, collecting longitudinal data on sexual behavior is essential for understanding and preventing HIV transmission.

ACKNOWLEDGMENTS

Author affiliations: Theoretical Biology and Biophysics Group, Los Alamos National Laboratory, Los Alamos, New Mexico (E. O. Romero-Severson, T. Leitner); Department of Infectious Disease Epidemiology, Imperial College London, London, United Kingdom (E. Volz); Department of Epidemiology, School of Public Health, University of Michigan, Ann Arbor, Michigan (J. S. Koopman); and Department of Statistics, College of Literature, Science and the Arts, University of Michigan, Ann Arbor, Michigan (E. L. Ionides).

This work was supported by the National Institutes of Health (grant R01AI78752 to J.S.K. and E.O.R.-S. and grant R01AI087520 to T.L. and E.O.R.-S.) and the National Science Foundation (grant DMS 1308919 to E.L.I.).

Conflict of interest: none declared.

REFERENCES

- Hall HI, Song R, Rhodes P, et al. Estimation of HIV incidence in the United States. *JAMA*. 2008;300(5):520–529.
- Wawer MJ, Gray RH, Sewankambo NK, et al. Rates of HIV-1 transmission per coital act, by stage of HIV-1 infection, in Rakai, Uganda. *J Infect Dis*. 2005;191(9):1403–1409.
- Pinkerton SD. Probability of HIV transmission during acute infection in Rakai, Uganda. *AIDS Behav*. 2008;12(5):677–684.
- Hollingsworth TD, Anderson RM, Fraser C. HIV-1 transmission, by stage of infection. *J Infect Dis*. 2008;198(5):687–693.
- Cohen MS, Shaw GM, McMichael AJ, et al. Acute HIV-1 infection. *N Engl J Med*. 2011;364(20):1943–1954.
- Vittinghoff E, Douglas J, Judson F, et al. Per-contact risk of human immunodeficiency virus transmission between male sexual partners. *Am J Epidemiol*. 1999;150(3):306–311.
- Anderson RM, May RM. *Infectious Diseases of Humans: Dynamics and Control*. New York, NY: Oxford University Press; 1992.
- Smith A, Miles I, Le B, et al. Prevalence and awareness of HIV infection among men who have sex with men—21 cities, United States, 2008. *MMWR Morb Mortal Wkly Rep*. 2010;59(37):1201–1207.
- Rapatski BL, Suppe F, Yorke JA. HIV epidemics driven by late disease stage transmission. *J Acquir Immune Defic Syndr*. 2005;38(3):241–253.
- Koopman JS, Jacquez JA, Welch GW, et al. The role of early HIV infection in the spread of HIV through populations. *J Acquir Immune Defic Syndr Hum Retrovirol*. 1997;14(3):249–258.
- Anderson RM, Medley GF, May RM, et al. A preliminary study of the transmission dynamics of the human immunodeficiency virus (HIV), the causative agent of AIDS. *IMA J Math Appl Med Biol*. 1986;3(4):229–263.
- Miller JC, Slim AC, Volz EM. Edge-based compartmental modelling for infectious disease spread. *J R Soc Interface*. 2012;9(70):890–906.
- Goodreau SM, Cassels S, Kasprzyk D, et al. Concurrent partnerships, acute infection and HIV epidemic dynamics among young adults in Zimbabwe. *AIDS Behav*. 2012;16(2):312–322.
- Kim J-H, Riolo RL, Koopman JS. HIV transmission by stage of infection and pattern of sexual partnerships. *Epidemiology*. 2010;21(5):676–684.
- Abu-Raddad LJ, Longini IM Jr. No HIV stage is dominant in driving the HIV epidemic in sub-Saharan Africa. *AIDS*. 2008;22(9):1055–1061.
- Hyman JM, Li J, Stanley EA. The differential infectivity and staged progression models for the transmission of HIV. *Math Biosci*. 1999;155(2):77–109.
- Alam SJ, Romero-Severson E, Kim J-H, et al. Dynamic sex roles among men who have sex with men and transmissions from primary HIV infection. *Epidemiology*. 2010;21(5):669–675.
- Zhang X, Zhong L, Romero-Severson E, et al. Episodic HIV risk behavior can greatly amplify HIV prevalence and the fraction of transmissions from acute HIV infection. *Stat Commun Infect Dis*. 2012;4(1):1041.
- Romero-Severson EO, Alam SJ, Volz E, et al. Acute-stage transmission of HIV: effect of volatile contact rates. *Epidemiology*. 2013;24(4):516–521.
- Balode D, Skar H, Mild M, et al. Phylogenetic analysis of the Latvian HIV-1 epidemic. *AIDS Res Hum Retroviruses*. 2012;28(8):928–932.
- Skar H, Sylvan S, Hansson H-B, et al. Multiple HIV-1 introductions into the Swedish intravenous drug user population. *Infect Genet Evol*. 2008;8(5):545–552.
- Ionides EL, Bretó C, King AA. Inference for nonlinear dynamical systems. *Proc Natl Acad Sci U S A*. 2006;103(49):18438–18443.
- Ionides EL, Bhadra A, Atchadé Y, et al. Iterated filtering. *Ann Stat*. 2011;39(3):1776–1802.
- Romero-Severson EO, Alam SJ, Volz EM, et al. Heterogeneity in number and type of sexual contacts in a gay urban cohort. *Stat Commun Infect Dis*. 2012;4(1):4.
- Keeling MJ, Rohani P. Stochastic dynamics. In: *Modeling Infectious Diseases in Humans and Animals*. Princeton, NJ: Princeton University Press; 2007:190–230.
- McCullagh P, Nelder JA. *Generalized Linear Models*. 2nd ed. London, United Kingdom: Chapman and Hall Ltd.; 1989.
- Bretó C, Ionides EL. Compound Markov counting processes and their applications to modeling infinitesimally over-dispersed systems. *Stoch Proc Appl*. 2011;121(11):2571–2591.
- Bretó C, He D, Ionides EL, et al. Time series analysis via mechanistic models. *Ann Appl Stat*. 2009;3(1):319–348.
- King AA, Ionides EL, Bretó C, et al. *Package 'pomp': Statistical Inference for Partially Observed Markov Processes [R software]*. Vienna, Austria: R Foundation for Statistical Computing; 2009. <http://cran.r-project.org/web/packages/pomp/pomp.pdf>. Accessed April 24, 2015.
- R Development Core Team. *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing; 2006.
- Arulampalam MS, Maskell S, Gordon N, et al. A tutorial on particle filters for online nonlinear, non-Gaussian Bayesian tracking. *IEEE Trans Signal Process*. 2002;50(2):174–188.
- Ionides EL. Maximum smoothed likelihood estimation. *Stat Sin*. 2005;15:1003–1014.
- Rice JA. *Mathematical Statistics and Data Analysis*. Belmont, CA: Wadsworth Publishing Company; 1988.
- Burnham KP, Anderson DR. *Model Selection and Inference: A Practical Information-Theoretic Approach*. 2nd ed. New York, NY: Springer-Verlag New York; 2002.
- Funk S, Salathé M, Jansen VAA. Modelling the influence of human behaviour on the spread of infectious diseases: a review. *J R Soc Interface*. 2010;7(50):1247–1256.
- Parsons JT, Schrimshaw EW, Wolitski RJ, et al. Sexual harm reduction practices of HIV-seropositive gay and bisexual men: serosorting, strategic positioning, and withdrawal before ejaculation. *AIDS*. 2005;19(suppl 1):S13–S25.