



Methods

A Bayesian approach to synthesize estimates of the size of hidden populations: the Anchored Multiplier

Paul D Wesson,^{1*} Ali Mirzazadeh² and Willi McFarland²

¹Center for AIDS Prevention Studies and ²Department of Epidemiology and Biostatistics, University of California San Francisco, San Francisco, CA, USA

*Corresponding author. Center for AIDS Prevention Studies, Division of Prevention Science, University of California San Francisco, 550 16th St, 3rd Floor, San Francisco, CA 94158, USA. E-mail: paul.wesson@ucsf.edu

Editorial decision 23 May 2018; Accepted 1 June 2018

Abstract

Background: The multiplier method is one of the most frequently used population size estimation (PSE) methods for key populations, yet estimates from this method are often inconsistent with each other, other PSE methods and local knowledge. We developed a novel Bayesian approach, the 'Anchored Multiplier', which synthesizes estimates from multipliers coupled to an a priori estimate to arrive at a single consensus estimate and credible range.

Methods: Data for size estimation were collected from three cross-sectional bio-behavioural surveillance studies of people who inject drugs (PWID) in San Francisco, CA, USA (2005, 2009 and 2012). We demonstrate the application of the Anchored Multiplier and a Variance Adjusted-Anchored Multiplier using PSE produced by multipliers in the three surveys and the literature for the USA. Size estimates were compared with estimates from other available PSE methods.

Results: Using the Anchored Multiplier, we estimated the PWID population made up 2.41% [95% credible interval (CI): 1.9–2.85] of the adult population in 2005, 2.1% (95% CI: 1.8–2.48) in 2009 and 2.3% (95% CI: 2.03–2.61) in 2012. The Variance Adjusted-Anchored Multiplier calculated similar point estimates, with wider 95% credible intervals. Credible intervals from both approaches were substantially narrower than from other standard PSE methods and, unlike other methods, indicated that the prevalence of PWID was stable over time.

Conclusions: The Anchored Multiplier is a promising new approach to size estimation, which generates a single estimate to inform programmatic strategies to counter the HIV epidemic, and provides a robust denominator to quantify the burden of disease for key populations.

Key words: Bayes Theorem, Delphi technique, HIV infections, population size, population surveillance, substance abuse, intravenous

Key Messages

- Population size estimates are fundamental to epidemiology; providing a denominator for the population at risk and justification for resource mobilization.
- The multiplier method is the most frequently used population size estimation method for key populations, and yet estimates from this method are often biased and discrepant.
- We developed a Bayesian model to synthesize estimates from multiple multiplier methods and mitigate bias, while incorporating local knowledge.
- Using this method, we show that the population of injection drug users in San Francisco, CA, has remained stable over time, contrary to previous claims of a growing population.

Background

Population size estimation (PSE), especially for key populations at high risk for HIV, is fundamental to disease surveillance and public health resource mobilization. Increasingly, the Centers for Disease Control and Prevention (CDC), the World Health Organization (WHO) and the Joint United Nations Programme on HIV/AIDS (UNAIDS) are prioritizing PSE of key populations in their support for national HIV control programmes.^{1,2} Notably, population size estimates are needed to gauge progress towards the UNAIDS 90–90–90 goals. Unless the denominator of populations at risk for HIV is known, we cannot determine whether 90% of people living with HIV have been diagnosed, 90% are on anti-retroviral treatment and 90% are virally suppressed.³ Unfortunately, estimating the size of key populations is challenging and fraught with many uncertainties in the absence of a gold standard.⁴

One of the most frequently used PSE methods is the multiplier method because it leverages programmatic data with a survey of the key population.^{5,6} This method requires two data sources: the ‘benchmark’ (n), which is a count of the number of people within the target population who received a service or unique object, and the ‘multiplier’ (p), which is the proportion of people from a representative sample of the target population who report receiving the service or unique object. Dividing the benchmark by the multiplier gives an estimate of the size of the target population (e).

$$\text{Multiplier Method} = e = \frac{n}{p}$$

The ease with which this PSE method is integrated into bio-behavioural surveillance surveys of key populations speaks to its popularity among investigators.⁷ In fact, surveys can simultaneously integrate multiple multipliers into one survey of a key population at little cost.

Despite being one of the most frequently used PSE methods, the multiplier method may be especially vulnerable to

biased estimates. A recent systematic review found that multiple estimates using the multiplier method (i.e. using different benchmarks) are often discrepant and the confidence intervals do not overlap.⁶ We developed and propose an improvement to the multiplier method that can recover from potentially biased data and reduce variability across multiple estimates. This novel method, which we call the ‘Anchored Multiplier’, uses a Bayesian framework to synthesize multiple data points, namely estimates from several multipliers, and stakeholder knowledge or estimates from the published literature as a priori belief. Estimates from the multiplier methods (likelihood) and prior belief (prior) are converted into probability distributions and combined according to Bayes Theorem to generate a posterior probability distribution that reflects the updated knowledge of the population size, driven by the strength (i.e. precision) of the individual inputs, as detailed below:

$$p(\theta|x) \propto p(x|\theta)p(\theta)$$

$$\text{posterior} \propto \text{likelihood} \times \text{prior}$$

With this approach, we arrive at a ‘consensus estimate’ that synthesizes both local or published knowledge and empirical data. Thus, estimates from the multiplier methods are ‘anchored’ to prior knowledge, limiting the influence of potentially biased estimates, reducing the influence of chance and making estimates more reasonable and acceptable to stakeholders. Here, we describe this novel PSE method and apply it as a case study to estimate the size of the population of people who inject drugs (PWID) in San Francisco, CA, using published data from 2005, 2009 and 2012.⁸

Methods

We used published data from the San Francisco Department of Public Health (SFDPH) to re-estimate the

number of PWID in San Francisco. SFPDPH conducted three cross-sectional bio-behavioural surveillance studies of PWID in 2005, 2009 and 2012 as part of the National HIV Behavioral Surveillance (NHBS).^{8–10} Each round of NHBS used respondent driven sampling (RDS), a peer referral-based method of recruitment, to sample and recruit study participants.¹¹ Full details of the study design and main outcomes have previously been reported.^{9,10} Chen *et al.*⁸ recently used data from these bio-behavioural surveillance studies to estimate the number of PWID in San Francisco using several PSE methods: namely, the multiplier method, the wisdom of the crowds and the successive sampling population size estimation method (SS-PSE). The median of all point estimates from each PSE method for each NHBS round was reported as the final estimate for the number of PWID in San Francisco. The same approach was taken for the upper and lower bounds of each estimate. We used the calculated population size estimates from each multiplier method reported by Chen *et al.* as input for our demonstration of the Anchored Multiplier.

The number of PWID in San Francisco were converted to population proportions using the number of adults living in San Francisco in 2005 as the denominator (777 660: U.S Census Bureau).¹² We then estimated the beta distribution that closely matched the details of the estimated population size; that is, the mean of the beta distribution matched the population proportion, and at least 70% of the probability density was contained within the upper and lower bounds of the original estimate (also converted to

population proportions). Beta distributions are probability distributions that are frequently used to describe prior uncertainty about disease prevalence. They are flexible and mathematically convenient for quantities constrained to lie between 0 and 1.¹³ We estimated the beta distributions for all priors and estimates from the multiplier methods. The population size estimates, population proportions and shape parameters for corresponding beta distributions for each of the multiplier method estimates are reported in Table 1.

We calculated the variance for each population size estimate, assuming a symmetrical distribution, using the following equation:

$$\left(\frac{\text{Upper bound} - \mu}{1.96}\right)^2 = \text{variance}$$

where μ refers to the population size estimate as a proportion. We then used the following equations to calculate the alpha (α) and beta (β) shape parameters for the beta distribution from the mean (μ) and variance of the estimated population size:^{14,15}

$$\alpha = -\frac{\mu(\sigma^2 + \mu^2 - \mu)}{\sigma^2}$$

$$\beta = \frac{(\sigma^2 + \mu^2 - \mu)(\mu - 1)}{\sigma^2}$$

Table 1. Population size estimates of people who inject drugs (PWID), population prevalence and beta distribution shape parameters from multiplier methods. San Francisco, CA 2005, 2009, 2012

Year	Source	Estimate	% Prevalence	Shape parameters	
				alpha	beta
2005	HIV/AIDS surveillance cases	23 779 (16 027 to 31 534)	3.06 (2.06 to 4.05)	35.55181	1126.272
	Arrests	20 909 (14 960 to 26 861)	1.3 (0.42 to 2.19)	8.084233	613.7798
	Participants of the UFO Study ³²	10 158 (3273 to 17 044)	2.69 (1.92 to 3.45)	46.91609	1697.177
	Walden House group home clients	10 130 (2998 to 17 263)	1.3 (0.38 to 2.22)	7.55863	573.8744
2009	San Francisco City Clinic HIV tests	81 500 (658 to 162 343)	10.5 (0.1 to 20.9)	3.39032	28.95956
	San Francisco City Clinic STI tests	7000 (2945 to 11 056)	0.9 (0.4 to 1.4)	11.33028	1247.399
	HIV/AIDS surveillance cases	45 315 (24 576 to 66 057)	5.8 (3.2 to 8.5)	17.20889	278.1164
	San Francisco General Hospital ER visits	18 250 (4675 to 31 826)	2.3 (0.6 to 4.1)	6.755769	281.1177
	Anonymous HIV tests	12 857 (0 to 41 780)	1.6 (0.0 to 5.4)	0.7300264	43.42587
	Walden House group home clients	5200 (1003 to 9398)	0.7 (0.1 to 1.2)	5.848229	868.7545
2012	The Stonewall Project applications	64 000 (37 951 to 90 050)	8.2 (4.9 to 11.6)	21.19706	236.3671
	The Stonewall Project treatment	60 000 (0 to 178 574)	7.7 (0.0 to 23.0)	0.8305923	9.934714
	HIV/AIDS surveillance cases	43 973 (20 943 to 67 006)	5.7 (2.7 to 8.6)	13.15351	219.4655
	Bayview-Hunter's Point methadone clients	34 600 (0 to 75 615)	4.5 (0.0 to 9.7)	2.567747	55.14422
	Davies Medical Center, overdose visits	11 250 (0 to 28 111)	1.5 (0.0 to 3.6)	1.671008	113.8379
	San Francisco City Clinic STI tests	8057 (3450 to 12 667)	1.0 (0.4 to 1.6)	11.60237	1108.255

ER, emergency room; STI, sexually transmitted infection.

Bayesian priors were selected from the literature. Lansky *et al.*¹⁶ reported 2.6% [95% confidence interval (CI): 1.8%–3.3%] of the adult US population had ever injected drugs in their lifetime, based on a meta-analysis of data from four national household-based surveys. This became our prior for the proportion of the population who had ever injected drugs. To assess the robustness of the Anchored Multiplier to the choice of priors, we specified two additional priors. Lansky *et al.* also reported 0.3% (95% CI: 0.19%–0.41%) of the US population had injected drugs in the past year. This became our prior for the proportion of the population who recently injected drugs. Finally, we specified a uniform prior distribution, representing little prior knowledge of the population size, with a lower bound of zero and an upper bound of 12.88%. This upper bound corresponds to the upper bound of the proportion of PWID in San Francisco in 2005, estimated by the SS-PSE method,^{17,18} the highest estimated upper bound among the PSE methods used by Chen *et al.*⁸

Analysis

Because the beta distribution is the conjugate prior for the binomial distribution, we used the binomial distribution as the sampling distribution to calculate the Anchored Multiplier. For each estimate of the population size from the multiplier method, the alpha shape parameter was used as the number of ‘successes’ and the sum of the alpha and beta shape parameters became the ‘sample size’ for the binomial distribution. We used the ‘rjags’ package in R statistical software to specify our Bayesian model and to use Markov Chain Monte Carlo (MCMC) simulation to estimate the posterior distribution.^{19,20} To calculate the Anchored Multiplier, the prior (beta) distribution was combined with the data (binomial) distribution for a single multiplier method estimate. MCMC diagnostic plots of the resulting posterior distribution (trace plots, autocorrelation plots, density plots) were visually inspected using the ‘bayesplot’ R package to assess model convergence, following the recommendations of Hamra *et al.*^{21,22} Shape parameters of the corresponding beta distribution were calculated from this posterior distribution, to be used as the new prior distribution, which was then combined with the data (binomial) distribution from another multiplier method estimate. This process continued iteratively until all multiplier method estimates were incorporated into the Bayesian model. For all simulations we specified three chains, each with 5000 iterations and a burn-in period of 2500. From the final posterior distribution, we calculated the mean, 2.5th percentile, and 97.5th percentile to report the point estimate and 95% credible interval for the Anchored Multiplier.

To account for the extremely narrow credible intervals resulting from increasingly strong priors that result from the incorporation of more data, we incorporated additional variance into the estimation of the posterior distributions. We followed the methodology for calculating random effects variance from the meta-analysis literature.²³ The calculations for this Variance Adjusted-Anchored Multiplier are detailed in [Supplementary materials](#), available as [Supplementary data](#) at *IJE* online.

For comparison, we plotted the estimates from the Anchored Multiplier and the Variance Adjusted-Anchored Multiplier with other population size estimates from Chen *et al.* and visually inspected these plots for trends over time by method. Linear trends were further assessed quantitatively by metaregression (‘metareg’ macro in Stata^{24,25}), regressing the estimated proportion of PWID on year. We subtracted 2005 from each year, and divided the difference by five so that the coefficient for year could be interpreted as the expected unit change in the population proportion for each 5-year interval. Last, as a separate sensitivity analysis, we assessed the impact of varying both the priors and the data on the final estimated population size. Following a jackknife approach, we estimated the population size using the Anchored Multiplier after leaving out one of the multipliers from the 2005 cross-sectional study. We repeated this analysis for all combinations of leaving one multiplier out, for the three priors used in our main study. To provide a comparison with standard practice, we again implemented the jackknife approach and reported the median of the remaining three multipliers, for all ‘leave one out’ combinations.

Results

The Anchored Multiplier method estimates the population of PWID to be 2.42% of the adult population of San Francisco in 2005, or roughly 18 820 adults (95% CI: 15 553–22 241). This estimation was based on a prior estimate that 2.6% of the adult population had ever injected drugs in their life, and data from four service multipliers applied to the San Francisco PWID population in 2005. When the between-method variance is incorporated (Variance Adjusted-Anchored Multiplier), the PWID population proportion increases to 2.47% (19 208 adults), and the 95% credible interval widens to reflect greater uncertainty (95% CI: 13 842–25 740). These results are depicted as a forest plot, comparing estimates from the Anchored Multiplier with the prior and the data inputs ([Figure 1](#)).

Population size estimates from individual multiplier methods ranged from 1.3% to 3.06% of the adult population in San Francisco in 2005, 0.7% to 10.5% in 2009 and 1% to 8.2% in 2012 ([Figure 2](#)). Final point estimates from




Figure 1. Population proportion of people who inject drugs (PWID) in San Francisco, CA (2005). *Anchored Multiplier-VA = Variance-Adjusted Anchored Multiplier.





Figure 2. Comparison of the proportion of people who inject drugs (PWID) in San Francisco, CA, at three cross-sections using different population size estimation (PSE) methods. *Anchored Multiplier-VA = Variance-Adjusted Anchored Multiplier.

Chen *et al.* suggest an increase in the population proportion of PWID over time; however, the upper and lower bounds in 2009 and 2012 are too wide to discern if this increasing trend is meaningful. Similarly, the 95% credible intervals for the SS-PSE in 2005 and 2009 both range from 0% to over 10% of the adult population in San Francisco. In comparison, the 95% credible intervals for both the Anchored Multiplier and the Variance Adjusted-Anchored

Multiplier are substantially narrower. The overlapping 95% credible intervals for each of the Anchored Multiplier methods indicate that this population proportion is stable over time. All 95% confidence intervals for coefficients for 'year' from the metaregression included the null, confirming no meaningful change in the population proportion from 2005 to 2012 (Supplementary Table 1, available as Supplementary data at *IJE* online).

Table 2. Anchored Multiplier estimates of the prevalence of people who inject drugs in San Francisco at three different cross-sections, using three different priors

	Prior = 2.6% (1.8% to 3.3%)		Prior = 0.3% (0.19% to 0.41%)		Prior = uniform (0, 12.88%)	
	Anchored Multiplier	Variance-Adjusted	Anchored Multiplier	Variance-Adjusted	Anchored Multiplier	Variance-Adjusted
2005	2.41 (1.9 to 2.85)	2.48 (1.77 to 3.32)	0.94 (0.79 to 1.11)	0.50 (0.38 to 0.64)	2.43 (1.99 to 2.93)	2.47 (1.60 to 3.54)
2009	2.1 (1.8 to 2.48)	2.81 (2.13 to 3.55)	1.06 (0.91 to 1.23)	0.67 (0.53 to 0.83)	2.11 (1.78 to 2.46)	2.92 (2.14 to 3.80)
2012	2.3 (2.03 to 2.61)	3.05 (2.36 to 3.82)	1.26 (1.10 to 1.43)	0.76 (0.60 to 0.93)	2.27 (1.96 to 2.60)	3.21 (2.43 to 4.11)

**Figure 3.** Sensitivity analysis of the estimated proportion of people who inject drugs (PWID) in San Francisco, CA, in 2005, comparing the Anchored Multiplier (with three different priors) with calculating the median of individual estimates. *VA = Variance-Adjusted estimate. **Legend and x-axis indicate which multiplier method estimate is left out of the data synthesis.

Using a uniform prior as a starting prior value, final estimates for both the Anchored Multiplier and the Variance Adjusted-Anchored Multiplier were similar to estimates when the prior for lifetime injection drug use was used (Table 2). The 95% credible intervals for both sets of Anchored Multiplier estimates derived from these two different prior values indicate that the estimates are not statistically different from each other. In contrast, when we used the prior that reflected recent injection drug use (0.3%; 0.19%–0.41%), estimates from the Anchored Multiplier and Variance Adjusted-Anchored Multiplier were substantially decreased. Population proportions ranged from 0.94% to 1.26% for the Anchored Multiplier and were lower when adjusting the variance. Although there was an increasing trend in point estimates, these estimates were not meaningfully different from each other, indicated by their marginally overlapping credible intervals.

The results of the sensitivity analysis are shown in Figure 3. Jackknife estimates using the Anchored

Multiplier were always in agreement with each other, in terms of similar point estimates and 95% credible intervals. In contrast, jackknife estimates using the standard approach of calculating the median of estimates showed greater variability in point estimates and marginally overlapping intervals.

Discussion

In this study, we developed a novel PSE method that uses a Bayesian framework to synthesize prior knowledge concerning the size of a population with empirical results using the multiplier method. In contrast to the previous PSE analysis that used the median of all estimates,⁸ our results suggest that the population proportion of people who have ever injected drugs in San Francisco has remained stable from 2005 to 2012, as indicated by the Variance Adjusted-Anchored Multiplier estimates (2005: 2.48%; 2009: 2.81%; 2012: 3.05%). Accounting for and incorporating

between-group (multiplier method) variance widened the credible intervals, reflecting greater uncertainty compared with the Anchored Multiplier estimates. These variance-adjusted credible intervals were still narrower than other reported size estimates of the same population, allowing for better inference of both population size and trends in the population size over time.

Iteratively incorporating more data resulted in narrower credible intervals ([Supplementary Figure 1a–e](#), available as [Supplementary data](#) at *IJE* online). In our application, the amount of data included through the iterative addition of estimates from the multiplier method resulted in 95% credible intervals that were extremely and perhaps implausibly narrow. We thought that these narrow credible intervals would communicate a false sense of certainty or raise suspicion in the value of using the resulting estimate; we therefore incorporated additional between-group variance as done in meta-analysis. In our case study, adjusting the variance did not meaningfully affect the point estimates (mean of the posterior distribution) when comparing estimates from the Anchored Multiplier with the Variance Adjusted-Anchored Multiplier, yet it did somewhat widen the credible intervals. Therefore, we recommend that investigators use and report the variance-adjusted estimate.

As is common in Bayesian analyses, the Anchored Multiplier is influenced by the strength of the prior. In real-world practice, the Anchored Multiplier would take into account confidence in peer-reviewed publications or stakeholder opinion on how reasonable the empirical estimates are while still producing estimates consistent with local evidence. In our case example, the mean estimates from the posterior distribution using recent injection drug use as a prior are substantially lower than mean estimates using lifetime injection drug use as a prior. This can be explained by the relative strengths of the priors. The width of the interval for the recent injection use prior is less than one-sixth of the width of the interval for lifetime injection drug use (0.22 vs 1.5). This difference in interval width translates to the prior on recent injection drug use being a stronger prior, relative to the prior on lifetime injection drug use, having greater influence on the posterior distribution of the population size.

In contrast, mean estimates from the posterior distribution when a uniform prior was used were similar to mean estimates when using the prior on lifetime injection drug use. Using a uniform prior essentially communicates little prior knowledge and an almost complete reliance on the data to estimate the population size. The fact that posterior estimates were similar when using a uniform prior and the lifetime injection drug use prior is likely due to the fact that the prior for lifetime injection drug use was similar to the multiplier method estimates from 2005 ([Figure 1](#)).

In fact, the point estimate for this prior fell between the point estimates for the multiplier methods, and the confidence interval for the prior overlapped the confidence intervals of the four estimates from the multiplier method. Because the prior and the data were in agreement in this case, replacing the informative prior with a uniform prior, thus relying heavily on information from the data (the multiplier methods), did not meaningfully change the posterior estimates.

The results of our sensitivity analysis demonstrate the robustness of estimates using the Anchored Multiplier, compared with estimates using the standard approach of calculating the median ([Figure 3](#)). Leaving out one of the multiplier method estimates did not meaningfully change the final Anchored Multiplier estimates. In contrast, the standard approach of calculating the median of estimates was clearly influenced by which estimate was removed from the calculation. This approach resulted in two disparate population size estimates (1.3% vs 2.69%) and marginally overlapping intervals. Thus, using a prior and this Bayesian framework appears to stabilize the final estimate.

The Anchored Multiplier relies on estimates from the multiplier method as input. Point estimates from the multiplier method may be biased due to strong dependence between participation in the service (benchmark) and the survey of the key population. Biases in the multiplier method arise from several other reasons, including: the service providing the benchmark count may provide the number of visits rather than the number of unique clients; people who are not a part of the target population (i.e. not PWID) may erroneously be included in the benchmark count; and study participants (contributing to the multiplier) may not remember that they visited the service in question. The Anchored Multiplier is not an automatic fix to these biases. However, the influence of these biases may be mitigated. If the estimates are biased, the anchored multiplier is unlikely to completely correct this unless the bias is known *a priori* and the prior is strong enough in the opposite direction. Furthermore, if the bias is known *a priori*, the investigator may widen the confidence intervals (effectively weakening the contributions of this data point to the posterior distribution).

Conclusion

As with most studies exploring population size estimation, the absence of a gold standard precludes our ability to definitively declare how much closer our method brings us to the truth. Unlike other studies, however, we propose a method that synthesizes available knowledge and data to reach a consensus estimate to inform public health action. Often estimates calculated from multiple implementations

of PSE methods do not agree with each other (i.e. point estimates differ and confidence intervals do not overlap).⁶ To resolve any disagreements and to put forward a single estimate, some researchers will simply take the median of all estimates. This approach assumes that all estimates are equally biased (or valid); albeit in different directions, such that the median of all estimates will balance out the biases. In other cases, estimates from PSE methods conflict with local stakeholder knowledge, jeopardizing the credibility and acceptance of these estimates. In practice, this conflict is often resolved using the modified Delphi approach, a subjective PSE method that presents local data to stakeholders, allowing them to update their previous beliefs about the population size now that they have seen the data (similar conceptually to the more formalized Bayesian approach we present in this paper). Again, a measure of central tendency (e.g. the median) summarizes the updated stakeholder beliefs. The Anchored Multiplier addresses both of these limitations in the current PSE field by: (i) combining information through an approach that effectively down-weights estimates where there is less certainty, rather than combine estimates in a way that assumes all estimates are equally valid and precise; and (ii) providing a framework to systematically synthesize multiple data points (including prior knowledge) in a way that is transparent. As such, the Anchored Multiplier is a promising new approach to estimating the size of key populations and thereby inform programmatic strategies to counter the HIV epidemic. A freely accessible web browser-based tool has been developed to implement the Anchored Multiplier. This online tool is available at the following link: [<http://globalhealthsciences.ucsf.edu/resources/tools>].

Supplementary Data

Supplementary data are available at *IJE* online.

Funding

This work was supported by the National Institutes of Health [T32 MH19105].

Conflict of interest: None declared.

References

- Schwartländer B, Stover J, Hallett T *et al.* Towards an improved investment approach for an effective response to HIV/AIDS. *Lancet* 2011;377:2031–41.
- Office of the US Global AIDS Coordinator. *PEPFAR 3.0—Controlling the Epidemic: Delivering on the Promise of an AIDS-Free Generation*. Washington, DC: Office of the US Global AIDS Coordinator, 2014.
- Joint United Nations Programme on HIV/AIDS. *90-90-90: An Ambitious Treatment Target to Help End the AIDS Epidemic*. Geneva: United Nations, 2014.
- UNAIDS/WHO Working Group on Global HIV/AIDS and STI Surveillance. *Guidelines on Estimating the Size of Populations Most at Risk to HIV*. Geneva: United Nations, 2010.
- Sabin K, Zhao J, Garcia Calleja JM *et al.* Availability and quality of size estimations of female sex workers, men who have sex with men, people who inject drugs and transgender women in low- and middle- income countries. *PLoS One* 2016;11: e0155150.
- Wesson P, Reingold A, McFarland W. Theoretical and empirical comparisons of methods to estimate the size of hard-to-reach populations: a systematic review. *AIDS Behav* 2017;21: 2188–2206.
- Johnston LG, Prybylski D, Raymond HF, Mirzazadeh A, Manopapiboon C, McFarland W. Incorporating the service multiplier method in respondent-driven sampling surveys to estimate the size of hidden and hard-to-reach populations: case studies from around the world. *Sex Transm Dis* 2013;40: 304–10.
- Chen Y, McFarland W, Raymond HF. Estimated number of people who inject drugs in San Francisco, 2005, 2009, and 2012. *AIDS Behav* 2016;20:2914–21.
- Evans JL, Hahn JA, Page-Shafer K *et al.* Gender differences in sexual and injection risk behavior among active young injection drug users in San Francisco (the UFO Study). *J Urban Health* 2003;80:137–46.
- Coffin PO, Jin H, Hurliaux E, Mirzazadeh A, Raymond HF. Trends in use of health care and HIV prevention services for persons who inject drugs in San Francisco: results from National HIV Behavioral Surveillance 2005–2012. *Drug Alcohol Depend* 2015;146:45–51.
- Heckathorn D. Respondent-driven sampling: a new approach to the study of hidden populations. *Soc Probl* 1997;44:174–99.
- United States Census Bureau. <https://www.census.gov/> (30 November 2016, date last accessed).
- Joseph L, Gyorkos TW, Coupal L. Bayesian estimation of disease prevalence and the parameters of diagnostic tests in the absence of a gold standard. *Am J Epidemiol* 1995;141: 263–72.
- Spiegelhalter DJ, Abrams KR, Myles JP. *Bayesian Approaches to Clinical Trials and Health-Care Evaluation*. Hoboken, NJ: Wiley, 2004.
- Cross Validated. *Calculating the Parameters of a Beta Distribution Using the Mean and Variance*. 2017. <https://stats.stackexchange.com/questions/12232/calculating-the-parameters-of-a-beta-distribution-using-the-mean-and-variance> (12 October 2016, date last accessed).
- Lansky A, Finlayson T, Johnson C *et al.* Estimating the number of persons who inject drugs in the United States by meta-analysis to calculate national rates of HIV and hepatitis C virus infections. *PLoS One* 2014;9:e97596.
- Handcock M, Gile K, Mar C. Estimating the size of populations at high risk for HIV using respondent-driven sampling data. *Biometrics*. 2015;71:258–66.

18. Wesson P, Handcock MS, McFarland W, Raymond HF. If you are not counted, you don't count: estimating the number of African-American men who have sex with men in San Francisco using a novel Bayesian Approach. *J Urban Health [Internet]* 2015;**92**:1052–64.
19. Plummer M. *rjags: Bayesian Graphical Models using MCMC*. R package version 4-6, 2016. <https://CRAN.R-project.org/package=rjags>.
20. R Core Team. *R: A Language and Environment For Statistical Computing*. Vienna: R Foundation for Statistical Computing, 2014.
21. Gabry J, Mahr T. *bayesplot: Plotting for Bayesian Models*. R package version 1.4.0, 2017. <https://CRAN.R-project.org/package=bayesplot>.
22. Hamra G, MacLehose R, Richardson D. Markov Chain Monte Carlo: an introduction for epidemiologists. *Int J Epidemiol* 2013;**42**:627–34.
23. Borenstein M, Hedges LV, Higgins JPT, Rothstein HR. *Introduction to Meta-Analysis*. Pondicherry, India: Wiley, 2009, pp. 72–73.
24. StataCorp. *Stata Statistical Software: Release 14*. College Station, TX: StataCorp LP, 2015.
25. Harbord R, Higgins J. *METAREG: Stata Module to Perform Meta-Analysis Regression*. Boston, MA: Boston College Department of Economics, 2004.