Opinion TRENDS in Microbiology Vol.14 No.12

Stochastic or deterministic: what is the effective population size of HIV-1?

Roger D. Kouyos[1], Christian L. Althaus[2] and Sebastian Bonhoeffer[1]

[1] Institute of Integrative Biology, ETH Zurich CHN, CH-8092 Zurich, Switzerland
[2] Theoretical Biology, Utrecht University, 3584 CH Utrecht, The Netherlands

Various studies have attempted to estimate the effective population size of HIV-1 to determine the strength of stochastic effects in within-host evolution. The largely discrepant estimates, the complexity of the concept of the effective population size and the resulting uncertainty about the underlying assumptions make the interpretation of these estimates difficult. Here, we explain the concept and critically assess the current estimates. We discuss the biologically relevant factors that affect the estimate and use of the effective population size. We argue that these factors lead to an underestimation of the effective population size and, thus, to an overestimation of the strength of stochastic effects in HIV-1 evolution.

Stochastic versus deterministic modeling of HIV-1

There is considerable debate in the literature as to whether the population dynamics and evolution of HIV-1 within a patient are more appropriately modeled as deterministic or stochastic processes [1–8]. Deterministic models assume an infinite population size and, thus, are only appropriate when population sizes are large. Given that there are at least 10^7–10^8 HIV-1-infected cells in a typical patient [9], this assumption seems to be justified at first sight. Importantly, however, the total (or census) population size is not the relevant quantity to determine whether processes can be described deterministically, as illustrated in the following thought experiment. For simplicity, assume that virus replication occurs in discrete generations. If all 10^8 infected cells in the current generation were infected by virus produced by a small subset of n randomly chosen cells in the previous generation, then this population of infected cells behaves as if it consisted of n cells. For a small n (e.g. $n = 10$ or 100), the population is obviously subject to stochastic effects, irrespective of its large census population size.

Assessing stochastic effects

Assessing the strength of stochastic effects has to be based on an 'effective' population size measure (N_e) that takes into account effects such as the one mentioned above. Several researchers have attempted to estimate N_e for HIV-1 [1–4,6–8,10]. However, diverging estimates and the inherent complexity of the concept of the effective population size have created considerable confusion as to how to interpret these estimates. Here, we will (i) provide an intuitive explanation of the concept of N_e; (ii) critically assess the current estimates; and (iii) discuss the applicability of these estimates. We argue that there is no universally applicable N_e for HIV-1. Using the current estimates of the effective population size of HIV-1 uncritically to assess the strength of stochastic effects can, therefore, be as misleading as trusting the apparently more naïve 'headcount'. Specifically, we argue that most current estimates of N_e overestimate the strength of stochastic effects for phenomena such as the evolution of drug resistance.

Effective population size

The difficulty in determining the appropriate population size has long been known to population geneticists. The concept of N_e was introduced to build a bridge between complex natural populations and their idealized mathematical representation and, thus, to enable an estimation of quantities of interest (which are hard to measure directly in natural populations) on the basis of the idealized model [11]. A graphical illustration of the concept of N_e is shown in Figure 1. In essence, the approach works as follows. First, one develops a stochastic model for the quantity of interest in which population size is a parameter. As long as the value for the population size is unknown, this idealized model is of little use. Thus, to estimate the relevant population size, one identifies another quantity (termed the calibration quantity) that can be determined easily in both the natural population and the model. N_e is then given as that population size for which the measured value of the calibration quantity in the idealized model is identical to that found in the natural population. Substituting the value of N_e for the population size in the model, one can then use the model to estimate other quantities of interest in natural populations. Hence, N_e provides the link between the model and the natural population by scaling the census population size to the appropriate size in the model.

Examples for such quantities of interest that are hard to measure directly in HIV-1-infected patients include the expected frequency of drug-resistant mutants in drug-naïve patients [5,12–15], the time of emergence of resistant mutants during therapy [12,13,16,17], the effect of recombination on the evolution of drug resistance [18] and the waiting times for the appearance of immune escape variants [19]. The calibration quantity is typically genetic diversity [1–4,6–8].

Corresponding author: Bonhoeffer, S. (seb@env.ethz.ch).
Available online 16 October 2006.

 0966-842X/$ – see front matter doi:10.1016/j.tim.2006.10.001

Glossary

Census population size: defined as the total number of individuals in a population.
Coalescent: a population genetic method that considers the ancestral history of a set of genes by developing a model describing how genetic lineages coalesce (i.e. merge) to common ancestry looking backwards in time.
Deterministic models: models that assume that the current state of a population fully determines the behavior of the population in the future. Deterministic models are often appropriate when populations are large.
N_e: the effective population size.
Neutral evolution: changes in allele frequencies by stochastic effects in the absence of fitness differences.
Stochastic models: models that assume that the future of a population is determined in part by its current state and in part by random effects. Random effects caused by sampling increase with decreasing population size and, therefore, stochastic models are often more appropriate than deterministic models when populations are small.
Wright–Fisher model (WFM): a model that is commonly used as the reference model to estimate N_e. The WFM describes discrete and non-overlapping generations in a population with a fixed size N [11,20]. Every generation, each of the N genomes undergoes mutation with a probability μ. Then the N genomes for the next generation are determined from the gene pool by drawing every offspring genome with uniform probability from the N parental genomes. Note that the WFM assumes selective neutrality.

Assumptions and limitations

The approach underlying N_e is appealing but, as always, the devil is in the details. There are important assumptions underlying the concept of N_e that represent potential pitfalls when applying this concept to real data. Conventionally, the idealized model is known as the Wright–Fisher model [11,20]. This model assumes that all individuals in the population have the same probability of reproduction and, thus, assumes neutrality. Consequently, when the Wright–Fisher model is used as the reference model, one can only use quantities that evolve neutrally both for calibration and for the measurement of interest. If this assumption is violated, misleading estimates can be obtained.

Although rarely done, it is also possible to use models other than Wright–Fisher as reference models. In particular, models can be used that incorporate selection to address quantities of interest that do not evolve neutrally. In any case, one needs to be aware that N_e is only defined in relation to a chosen calibration quantity [21] and, thus, implicitly reflects a particular evolutionary process (such as neutral or directional evolution). Care must be taken that the quantity of interest depends on the same evolutionary process as the calibration quantity, and that the model incorporates this evolutionary process.

Why is it important that the quantity of interest and the calibration quantity are subject to the same evolutionary process? Many factors, such as variation in reproductive success, selective sweeps and fluctuations in population size, are known to lead to N_e being smaller than the census population size in HIV-1 (Box 1). The difficulties in interpreting the estimates of N_e for intra-patient HIV-1 evolution often arise because these factors tend to have a different effect on the quantities of interest and of calibration. To illustrate these difficulties, we describe, in Box 2, a simulation that demonstrates how bottlenecks can have different effects on diversity at selected and at neutral sites.

Measurements of N_e in HIV-1

With the knowledge of the factors that influence the estimates of N_e and the difficulties that arise when comparing diversity at neutral sites with diversity at selected sites, the different attempts to measure N_e in HIV-1 can be assessed (Box 3). These attempts fall into two categories: those based on models of neutral evolution [1–3,6–8] and those based on models with selection [4]. Studies on the basis of models of neutral evolution have yielded small estimates of N_e that range from several hundreds to several thousands. By contrast, the study on the basis of a model with selection estimated that N_e is greater than 10^5–10^6.

Measurements assuming neutral evolution

As mentioned above, when models of neutral evolution are used as a reference, both the quantities of interest and of calibration need to evolve neutrally. Whether this is ful-

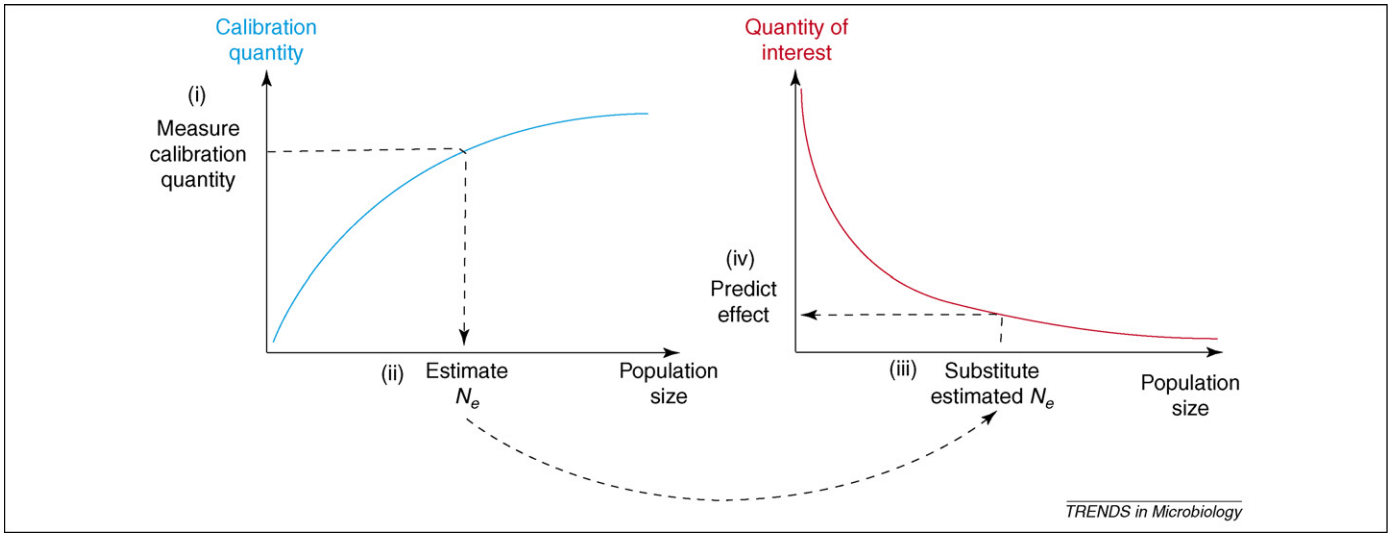

Figure 1. An illustration of the concept and use of N_e. The estimation of N_e is based on an idealized mathematical model in which population size is a parameter. To estimate a quantity of interest that is difficult to measure in a natural population, the approach taken is as follows: (i) One decides on an easily measurable quantity (calibration quantity) such as genetic diversity. (ii) N_e is defined as that population size for which the calibration quantity in the idealized model has the same value as that measured in the natural population. (iii) N_e is then substituted into the model to predict the value of the quantity of interest in the natural population (iv).

Box 1. Factors affecting the estimate of N_e

Examples of biological processes that lead to N_e being smaller than the census population size for neutral loci in HIV-1:

(i) **Variance in progeny number:** In the Wright–Fisher model, every progeny genome is chosen randomly from the parental genomes. In natural populations, however, often a few individuals produce the majority of the progeny, thereby increasing stochastic effects and decreasing estimates of N_e. Such a variance in progeny number can be a result of a heterogeneous environment or differences in the genetic background. For HIV-1, it is known that many infected cells (such as latently or defectively infected cells) produce little or no virus, resulting in a variance in reproductive output.

(ii) **Bottlenecks:** Temporary decreases in population size can greatly reduce genetic diversity and, therefore, reduce the corresponding estimates of N_e. Assuming neutral evolution, Wright [33] showed that for a periodically changing population size, N_e is given by the harmonic mean of the population sizes over all generations. In particular, this implies that even short periods of small population size have a disproportionately strong influence on N_e. HIV-1 populations periodically encounter bottlenecks either at transmission or through selective sweeps.

(iii) **Metapopulations:** Frost *et al.* [34] stated that the HIV-1 population within a host is described better by a collection of local subpopulations than by one well-mixed population. These subpopulations undergo frequent extinction and recolonization. Because every recolonization event functions as a 'local bottleneck' (a small number of viruses are founders of the subpopulation) this process again leads to a smaller N_e.

Importantly, these three types of factors have a different influence on stochastic effects and, thus, lead to different estimates of N_e, depending on whether one considers neutral or selected loci and depending on the population process studied.

filled is questionable. Regarding the calibration quantity, several studies [1–3,6,8] used polymorphisms of the HIV-1 *env* gene. Assuming neutral evolution of the env protein is problematic because the protein mediates viral entry and is targeted by the immune response. Moreover, the gene has a high ratio of nonsynonymous to synonymous nucleotide substitutions [22–24], which is indicative of selection. To justify a coalescent-based approach, which is based on neutral evolution, several of these studies [1,8] used statistical tests for selective neutrality. Indeed, these tests did not show a significant departure from neutrality. However, absence of evidence is not evidence of absence. A failure of these tests to detect selection could also be a result of limited statistical power. In another study using *gag–pol* sequences to estimate N_e [7], the authors state that these sequences are under a regime of strong purifying selection with ratios of nonsynonymous to synonymous nucleotide substitutions ranging between 0.01and 0.05.

However, we do not regard the possibility of selection acting on the calibration quantity as the main difficulty of this approach. Indeed, an analysis on the basis of third codon positions, which are expected to be under very weak selection, could lead to similar estimates provided the diversity at synonymous and nonsynonymous sites is not very different. The actual problem arises when applying these estimates of N_e to quantities of interest that are known to be under selection pressure (such as drug-resistance-associated polymorphisms) because factors such as those listed in Box 1 have different effects on neutral and selected sites. For example, the simulations in Box 2 show that if N_e is estimated on the basis of diversity at neutral sites and applied to diversity at selected sites, then bottlenecks can lead to a severe underestimation of N_e because they affect diversity at neutral sites more than at selected sites. In HIV-1, such bottlenecks occur either at transmission or as a consequence of selective sweeps caused by selection on linked loci [2]. Thus, bottlenecks might be one important factor that could lead to small estimates of N_e in HIV-1 [2,25].

Measurements allowing for selection

Rouzine and Coffin [4] use a model with selection to estimate N_e and thereby circumvent the problems associated with the assumption of selective neutrality. Shriner *et al.* [8] criticized this study in two ways. The first criticism is that Rouzine and Coffin overestimated N_e because they only considered polymorphic loci in their data. However, because they imposed the same restriction on their simulations, we do not agree with this criticism. The second criticism concerns the assumption of the constancy of the mutation rate across sites, which might indeed be problematic. A problem not mentioned in Shriner *et al.* [8] is that Rouzine and Coffin assumed such a low recombination rate that the effect of recombination on the estimate of N_e can be neglected. However, more recent estimates of the recombination rate [26] and the frequency of multiply infected cells [27] suggest that the effective recombination rate might be several orders of magnitude greater that than that used by Rouzine and Coffin. A high recombination rate could explain the simultaneous presence of all allele combinations even in small populations (Box 3) and consequently would lead to a lower (and hence less informative) estimate for the lower bound of N_e.

Concluding remarks and future perspectives

Despite the contradictory estimates of N_e, a valuable contribution of the described studies is that they highlight that the large census population size of HIV-1 does not imply the absence of stochastic effects in within-host HIV-1 evolution. For many quantities of interest, N_e might indeed be considerably smaller than the census population size. In fact, it is possible that N_e for the diversity at neutral sites is as small as 1000. Nevertheless, an N_e of 1000 cannot simply be plugged into models. For example, consider the pre-existence of drug-resistant mutants in a drug-naïve patient. One way to address whether such mutants are present when the patient is put on therapy is to determine the probability of the recent production of resistant mutants. Here, the census population size is the relevant quantity for the following reason. Within a time span of one to two days, much of the infected cell population is replaced [28,29]. This implies that in the order of 10^7–10^8 cells get infected in this time frame. Because mutations occur after the infection of a cell at a rate of 3.4×10^{-5} per nucleotide [30], ~100–1000 of the newly infected cells are expected to carry any single point mutation. The appearance of such mutations can, therefore, be

Box 2. Differential impact of bottlenecks on neutral and selected diversity

Here, we consider an extension of the Wright–Fisher model incorporating selection and bottlenecks. Between mutation and sampling steps, the frequency of the mutant allele is increased by a factor $1 + s$ relative to the wild type. The selection coefficient s is either 0 (neutral evolution) or -0.1 (selection against the mutant allele). In addition, the population can go through bottlenecks. Bottlenecks are implemented as follows. Every 1000 generations, ten individuals are sampled from the population. The simulations are then continued for four generations with this reduced population size. Thereafter, the population size is restored to its original size. We simulate populations with 10^6 individuals with and without bottlenecks, and populations with 2500 individuals in absence of bottlenecks. (Note that 2500 corresponds to the harmonic mean of the population size over time used in the simulations with bottlenecks; see Box 1).

Bottlenecks strongly reduce neutral diversity. With bottlenecks, for most of the time the allele frequency is close to 1 or 0, whereas it is always close to 0.5 without bottlenecks (Figure Ia,c). As expected, a constant population with 2500 individuals produces a similar pattern to the one obtained with a large population size and periodic bottlenecks. Thus, genetic variation at neutral loci suggests that the bottlenecks can be taken into account by reducing the population by a factor of 40.

The picture is completely different for selected loci. In the simulations bottlenecks have little effect on diversity (Figure Ib,d), which suggests that N_e is close to the census population size. Note that at least one copy of the mutant allele is present most of the time. By contrast, a population of a constant but reduced size $N_e = 2500$ (scaling the population size on the basis of diversity at neutral loci) shows a very different frequency distribution. In this case, the mutant allele is absent most of the time.

The reason for this discrepancy is that, in the absence of bottlenecks, the expected diversity at neutral loci considerably exceeds that at selected loci. After a reduction through a bottleneck, it takes longer for the neutral diversity to recover to its expected level. Because high diversity results in high estimates of N_e, using neutral diversity as a calibration quantity leads to an underestimation of N_e for selected loci in the presence of bottlenecks.

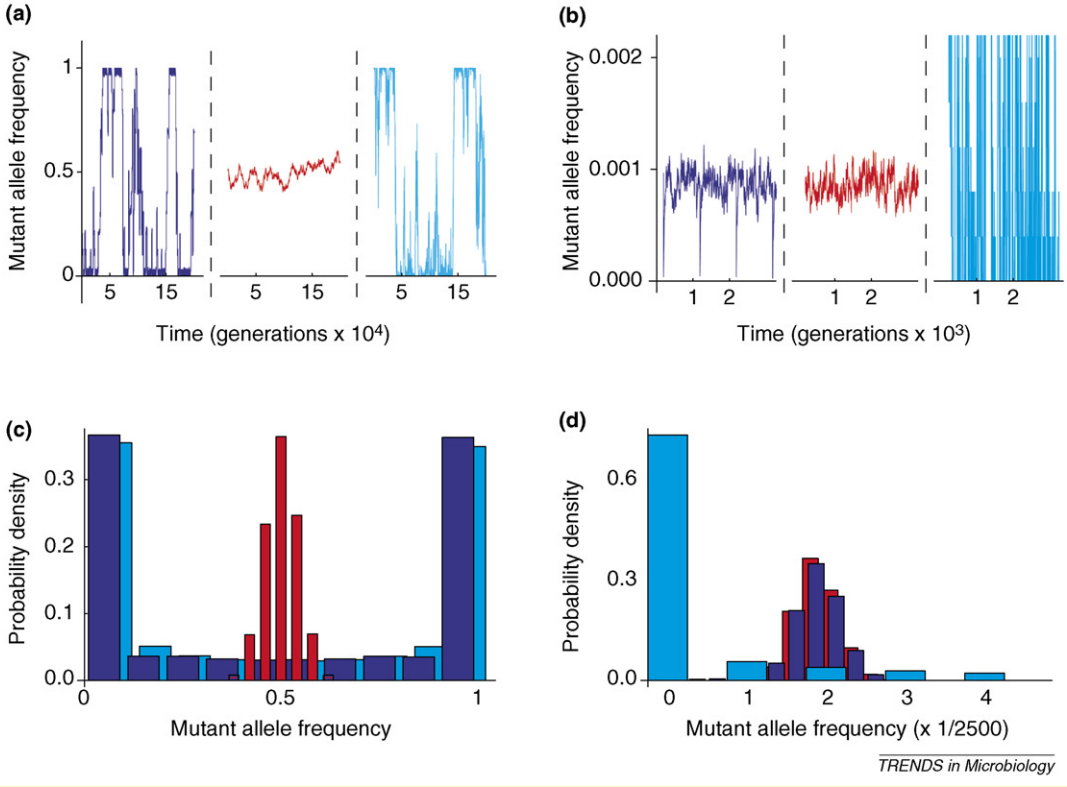

Figure I. A time-course of the frequency of the mutant allele and the corresponding probability density histograms describing the fraction of time the mutant allele has a particular frequency during a sufficiently long simulation. In **(a)** and **(c)**, the wild type and the mutant allele have the same fitness. In **(b)** and **(d)**, the mutant allele is selected against with a selection coefficient $s = -0.01$. Dark blue represents a large, bottlenecked population (10^6 individuals; bottlenecks: ten individuals), red represents a constantly large population (10^6 individuals), and light blue represents small populations (2500 individuals). The mutation rate is $\mu = 3.5 \times 10^{-5}$. The scales of the axes differ between plots; in plots (c) and (d), the histograms are shifted slightly along the x axis to enhance visibility.

described deterministically. Using an N_e of 1000 wrongly suggests that such a mutant would occur only about once every one to two months. This general problem is not confined to the example above but also appears, for example, when applying estimates of N_e to assess the effect of recombination on the evolution of drug resistance [18].

The limitations of the studies discussed here highlight the necessity to choose calibration quantities that adequately reflect the behavior of the quantity of interest. To adapt the approach of Leigh Brown [1] based on standing genetic diversity would require developing reference models that incorporate selection. Alternatively, Rouzine and Coffin's approach [4] could be extended to multiple linked loci. Such an approach would yield a more informative lower bound for N_e because evolution at multiple loci shows stochastic effects even for very large population sizes [31,32]. In both cases, this would create substantial experimental challenges because it is necessary to obtain better data, for example, for genetic diversity at selected sites and the selection coefficients of the corresponding alleles. Because increasing efforts need to

Box 3. Estimates of effective population sizes in HIV-1

Estimating N_e on the basis of models of neutral evolution:
Leigh Brown [1] estimated an N_e of ~1000 by applying the following method [35] to *env* polymorphisms: The probability that a given gene tree has produced the sequences [36] is multiplied by the probability that this tree is produced by a population of size N (the last probability being calculated with the coalescent). Summing these products for a representative sample of trees yields the total probability that a population of size N has produced the sequences. Finally, the population size is estimated as that value of N yielding the largest total probability. Importantly, the coalescent assumes neutrality. To justify this assumption, the author used Tajima's D test [37], which yielded no significant deviation from the (neutral) null hypothesis. Several subsequent studies reproduced this result using similar approaches [2,3,6]. Of note, Shriner *et al.* [8] estimated N_e allowing for recombination and obtained results consistent with the other studies.

Achaz *et al.* [7] estimated N_e to be between 10^3 and 10^4 by applying a test for population subdivision [38] both on sequence samples (*gag–pol* region) from different time points and on simulation data from a coalescent process. The time difference for which the test can distinguish between two samples is determined for both datasets. For neutral loci, this quantity scales with the population size. Thus, the ratio of population sizes is the ratio of the respective separation times.

Estimating N_e on the basis of a model with selection:
Rouzine and Coffin developed an alternative approach [4] on the basis of a two-locus–two-allele model with selection, which yields a lower limit for N_e based on the observation that such systems usually lack at least one of the four combinations of alleles for $N_e < 10^5$. However, HIV-1 sequence data showed that for a large fraction of the pairs of polymorphic loci all combinations were present, suggesting that N_e is larger than 10^5 (the most likely estimate being 10^6).

go into the estimation of the calibration quantity, naturally the question arises as to whether one could not determine the quantity of interest directly with similar effort.

Acknowledgements

We thank Viktor Müller, Lucy Crooks, Roland Regoes, Olin Silander, Marcel Salathé, Rob de Boer and Joachim Hermisson for valuable comments, and the Swiss National Science Foundation for financial support.

References

1 Leigh Brown, A.J. (1997) Analysis of HIV-1 env gene sequences reveals evidence for a low effective number in the viral population. *Proc. Natl. Acad. Sci. U. S. A.* 94, 1862–1865
2 Nijhuis, M. *et al.* (1998) Stochastic processes strongly influence HIV-1 evolution during suboptimal protease-inhibitor therapy. *Proc. Natl. Acad. Sci. U. S. A.* 95, 14441–14446
3 Rodrigo, A.G. *et al.* (1999) Coalescent estimates of HIV-1 generation time *in vivo*. *Proc. Natl. Acad. Sci. U. S. A.* 96, 2187–2191
4 Rouzine, I.M. and Coffin, J.M. (1999) Linkage disequilibrium test implies a large effective population number for HIV *in vivo*. *Proc. Natl. Acad. Sci. U. S. A.* 96, 10758–10763
5 Rouzine, I.M. *et al.* (2001) Transition between stochastic evolution and deterministic evolution in the presence of selection: general theory and application to virology. *Microbiol. Mol. Biol. Rev.* 65, 151–185
6 Seo, T.K. *et al.* (2002) Estimation of effective population size of HIV-1 within a host: a pseudomaximum-likelihood approach. *Genetics* 160, 1283–1293
7 Achaz, G. *et al.* (2004) A robust measure of HIV-1 population turnover within chronically infected individuals. *Mol. Biol. Evol.* 21, 1902–1912
8 Shriner, D. *et al.* (2004) Influence of random genetic drift on human immunodeficiency virus type 1 *env* evolution during chronic infection. *Genetics* 166, 1155–1164
9 Haase, A.T. *et al.* (1996) Quantitative image analysis of HIV-1 infection in lymphoid tissue. *Science* 274, 985–989
10 Frost, S.D. *et al.* (2000) Evolution of lamivudine resistance in human immunodeficiency virus type 1-infected individuals: the relative roles of drift and selection. *J. Virol.* 74, 6262–6268
11 Wright, S. (1931) Evolution in Mendelian populations. *Genetics* 52, 97–159
12 Bonhoeffer, S. *et al.* (1997) Virus dynamics and drug therapy. *Proc. Natl. Acad. Sci. U. S. A.* 94, 6971–6976
13 Bonhoeffer, S. and Nowak, M.A. (1997) Pre-existence and emergence of drug resistance in HIV-1 infection. *Proc. Biol. Sci.* 264, 631–637
14 Ribeiro, R.M. and Bonhoeffer, S. (2000) Production of resistant HIV mutants during antiretroviral therapy. *Proc. Natl. Acad. Sci. U. S. A.* 97, 7681–7686
15 Ribeiro, R.M. *et al.* (1998) The frequency of resistant mutant virus before antiviral therapy. *AIDS* 12, 461–465
16 Ribeiro, R.M. and Bonhoeffer, S. (1999) A stochastic model for primary HIV infection: optimal timing of therapy. *AIDS* 13, 351–357
17 Bocharov, G. *et al.* (2005) A genetic-algorithm approach to simulating human immunodeficiency virus evolution reveals the strong impact of multiply infected cells and recombination. *J. Gen. Virol.* 86, 3109–3118
18 Althaus, C.L. and Bonhoeffer, S. (2005) Stochastic interplay between mutation and recombination during the acquisition of drug resistance mutations in human immunodeficiency virus type 1. *J. Virol.* 79, 13572–13578
19 Liu, Y. *et al.* (2006) Waiting times for the appearance of cytotoxic T-lymphocyte escape mutants in chronic HIV-1 infection. *Virology* 347, 140–146
20 Fisher, R.A. (1930) The Genetical Theory of Natural Selection. *Clarendon*
21 Crow, J.F. (1954) Breeding structure of populations. II effective population number. In *Statistics and Mathematics in Biology* (Kempthorne, O., Bancroft, T., Gowen, J. and Lush, J., eds), pp. 543–556, Iowa State University Press
22 Bonhoeffer, S. *et al.* (1995) Causes of HIV diversity. *Nature* 376, 125
23 Seibert, S.A. *et al.* (1995) Natural selection on the *gag*, *pol*, and *env* genes of human immunodeficiency virus 1 (HIV-1). *Mol. Biol. Evol.* 12, 803–813
24 Yamaguchi, Y. and Gojobori, T. (1997) Evolutionary mechanisms and population dynamics of the third variable envelope region of HIV within single hosts. *Proc. Natl. Acad. Sci. U. S. A.* 94, 1264–1269
25 Leigh Brown, A.J. and Richman, D.D. (1997) HIV-1: gambling on the evolution of drug resistance? *Nat. Med.* 3, 268–271
26 Levy, D.N. *et al.* (2004) Dynamics of HIV-1 recombination in its natural target cells. *Proc. Natl. Acad. Sci. U. S. A.* 101, 4204–4209
27 Jung, A. *et al.* (2002) Multiply infected spleen cells in HIV patients. *Nature* 418, 144
28 Ho, D.D. *et al.* (1995) Rapid turnover of plasma virions and CD4 lymphocytes in HIV-1 infection. *Nature* 373, 123–126
29 Wei, X. *et al.* (1995) Viral dynamics in human immunodeficiency virus type 1 infection. *Nature* 373, 117–122
30 Mansky, L.M. and Temin, H.M. (1995) Lower *in vivo* mutation rate of human immunodeficiency virus type 1 than that predicted from the fidelity of purified reverse transcriptase. *J. Virol.* 69, 5087–5094
31 Iles, M.M. *et al.* (2003) Recombination can evolve in large finite populations given selection on sufficient loci. *Genetics* 165, 2249–2258
32 Rouzine, I.M. *et al.* (2003) The solitary wave of asexual evolution. *Proc. Natl. Acad. Sci. U. S. A.* 100, 587–592
33 Wright, S. (1938) Size of population and breeding structure in relation to evolution. *Science* 87, 430–431
34 Frost, S.D. *et al.* (2001) Genetic drift and within-host metapopulation dynamics of HIV-1 infection. *Proc. Natl. Acad. Sci. U. S. A.* 98, 6975–6980
35 Kuhner, M.K. *et al.* (1995) Estimating effective population size and mutation rate from sequence data using Metropolis-Hastings sampling. *Genetics* 140, 1421–1430
36 Felsenstein, J. (1981) Evolutionary trees from DNA sequences: a maximum likelihood approach. *J. Mol. Evol.* 17, 368–376
37 Tajima, F. (1989) Statistical method for testing the neutral mutation hypothesis by DNA polymorphism. *Genetics* 123, 585–595
38 Hudson, R.R. *et al.* (1992) A statistical test for detecting geographic subdivision. *Mol. Biol. Evol.* 9, 138–151