3-20-2003

Estimating the accuracy of polymerase chain reaction-based tests using endpoint dilution

Jim Hughes

University of Washington, jphughes@u.washington.edu

Patricia Totten

University of Washington, patotten@u.washington.edu

1 Introduction

Polymerase chain reaction (PCR)-based tests are commonly used to detect various infections (e.g. *Chlamydia trachomatis*, *Mycoplasma genitalium*, Human papilloma virus). Typically, such tests consist of a PCR amplification step to increase (by several orders of magnitude) the number of copies of a target DNA sequence in the sample, followed by detection of these "PCR products" by hybridization to a internal sequence after immobilization of the products on microwell plates or nytran filters. The result of the detection step is simply positive or negative for the target DNA sequence. Although such tests are universally acknowledged to be highly "sensitive" for the target DNA sequence - claims are often made that the test can detect a single molecule of the target DNA in a sample - no universally accepted definition of, or method for estimating, sensitivity exists. Some authors define sensitivity as the probability of a positive test when the target DNA sequence is present and specificity as the probability of a negative test when the target is absent (Rodrigo et al., 1997). However, this ignores the possibility that the probability of a positive test may vary depending on whether 1 copy or 20 copies of the target are present. Others characterize sensitivity in terms of the minimum number of molecules of target DNA that are detectable (e.g. Sykes et al, 1992) but this implies that the probability of detection is 1.0 at or above the detection limit and zero below it, which may not be true. Thus, it is unclear how these values should be interpreted.

In addition to the lack of a standard definition of sensitivity, there has been little work on estimating the sensitivity of a given test. In some cases, the test is implicitly assumed to be perfect - sensitivity and specificity equal to 1.0 (e.g. Shapiro, 1999). Rodrigo et al. (1997) explicitly incorporated a user-specified specificity and sensitivity into an analysis of HIV viral load levels. As noted above, however, they assumed, that the sensitivity was constant for all concentrations greater than 0. In addition, Rodrigo et al. did not indicate how the sensitivity could be estimated. Sykes et al. (1992) used an endpoint dilution experiment (see

2

below) to estimate an index of sensitivity they called the "mean number of copies detected". The authors treat this quantity like a detection limit although no explicit interpretation is given.

We propose that test sensitivity should be defined as the probability of a positive test and expressed as a function of the number of target DNA molecules in the sample (or specificity when the number of target molecules is zero). An endpoint dilution experiment may then be used to estimate this "sensitivity curve". Specifically, a sample containing a known concentration of the target DNA is serially diluted to yield solutions with varying (but known) concentrations of the target. Denote the concentration of target DNA molecules in the i'th dilution as λ_i, $i = 1...M$, ($\lambda_i = 0$ for the negative controls). At each of the M dilutions, k_i aliquots are sampled, amplified by PCR and tested for the presence of the target DNA. The observed data consist of the counts of positive (x_i) and negative tests $(k_i - x_i)$ at each dilution. This design is similar to a limiting dilution assay (LDA), which is often used to estimate an unknown concentration or proportion of targets in the original sample.

A simple estimate of specificity from such an experiment is the proportion of negative tests seen in the negative control samples. False positive tests in the negative controls may result if the PCR primers bind to and amplify non-target DNA sequences. Estimation of sensitivity as a function of the number of copies of the target DNA in the sample is more problematic, however. An important limitation of the endpoint dilution procedure described above is that, although the expected or average concentration of pathogen at the i'th dilution level is known, the researcher typically cannot control the exact number of DNA copies in each aliquot. Thus, false negative tests may occur for one of two reasons: 1) although the aliquot used for testing was taken from a solution with $\lambda_i > 0$, the aliquot contained no copies of the target DNA, or 2) although the aliquot used for testing contained one or more copies of the target DNA, the PCR procedure failed to amplify the DNA. It is the second reason that characterizes the true sensitivity of the test procedure. The first reason is simply an artifact of the experimental setup and is not related to the actual performance of the test. For this

reason we distinguish between the "observed" test sensitivity (the probability of detecting the pathogen in a sample taken from a solution with a known mean DNA concentration) and the "true" test sensitivity (the probability of detecting the pathogen in a sample containing a known number of the target DNA molecules).

Observed sensitivity may be estimated in a straightforward manner from the observed data using standard statistical tools such as logit or probit regression. In contrast, it is not possible to directly measure the true sensitivity since the exact number of DNA molecules in each sample is unknown. Nonetheless, it is possible, under reasonable assumptions, to estimate the true sensitivity and this gives a better indication of the performance of the test. For instance, suppose a 63% detection rate is observed when the mean concentration is 1.0 molecule of DNA per aliquot (i.e. observed sensitivity = 0.63). It turns out (details below) that this is the expected observed sensitivity when the true sensitivity is 1.0. Thus, it would be a waste of time to try to improve the test further, a conclusion that is not immediately obvious otherwise.

In section 2 we introduce a model for relating true sensitivity to the observed sensitivity and show how the single hit Poisson model described by Taswell (1981) may be used to estimate the true sensitivity under a particular set of assumptions about the form of the sensitivity curve. We then investigate ways of relaxing these assumptions. In section 3, we apply these methods to an endpoint dilution experiment designed to measure the sensitivity of a test for *M. genitalium*. We summarize our findings and discuss extensions in section 4.

2   Models

We define the true sensitivity and specificity in terms of a function $f(n)$ where $f(n) = \Pr(\text{positive test} \mid \text{exactly } n \text{ DNA molecules in the sample})$. Thus, specificity is equal to $1 - f(0)$ and sensitivity is defined by the function $f(n)$ for $n > 1$.

As noted above, f cannot be estimated from the data without further assumptions. The

4

key assumption is that, conditional on the mean concentration, λ, the actual number of DNA molecules, n, in an aliquot of volume v follows a Poisson distribution with mean $\mu = \lambda v$. That is,

$$Pr(n \mid \mu) = \frac{\mu^n e^{-\mu}}{n!}. \tag{1}$$

This assumption has been used extensively in the analysis of data from limiting dilution experiments (Sykes et al., 1992; Sharrock et al., 1990; Shapiro, 1999) and is reasonable if the dilution from which the aliquot is drawn is well mixed and there is no clumping of DNA molecules.

The observed sensitivity and specificity of the PCR procedure can then be characterized in terms of f and μ. Let $h(\mu)$ denote the observed sensitivity when the mean number of DNA molecules in the sample is known to be μ ($1 - h(\mu)$ is the specificity when $\mu = 0$):

$$h(\mu) = \sum_{n=0}^{\infty} \frac{f(n)\mu^n e^{-\mu}}{n!} \tag{2}$$

Note that even with perfect true sensitivity and specificity ($f(0) = 0$ and $f(n) = 1$ for $n > 1$), the observed sensitivity, h, is only 0.63 when $\lambda = 1$.

The functional form of f is unknown and in the following we investigate parametric, semiparametric (spline-based) and nonparametric models for f. Each has advantages and disadvantages. In each case, however, we restrict our attention to functions that are monotonically increasing in n since it seems reasonable to assume that the sensitivity must increase as the number of DNA molecules in the sample increases.

2.1 Parametric model for f

A number of parametric models for f are possible. We investigate the following,

$$f(n) = 1 - \phi(1 - \theta)^n. \tag{3}$$

This model may be interpreted mechanistically as follows: suppose that each of the n DNA molecules in the sample has an independent probability of $1 - \theta$ of "escaping" amplification

5

by PCR. If one or more of the n molecules are amplified then the sample will test positive while if none of the n molecules are amplified then the solution will test negative. In addition, there is a $1 - \phi$ probability that the sample will test positive even if no DNA molecules are present (so ϕ is the specificity). For this model, $h(\mu)$ in (2) reduces to the closed form

$$h(\mu) = 1 - \phi e^{-\theta \mu} \tag{4}$$

For $\phi = 1$ this model is identical to the single-hit Poisson model used by Taswell (1981) and others to analyze limiting dilution assay (LDA) data. In LDA experiments, however, θ is the fraction of a particular cell type in a mixture of cell types. Here, θ characterizes the test sensitivity. If $\phi = 1$, θ corresponds to the test sensitivity when a single DNA molecule is present and is equal to the inverse of the "mean number of copies detected" defined by Sykes et al. (1992). Use of the single hit Poisson model for the analysis of LDA data has been extensively studied (e.g. Taswell (1981), Cyr et al. (1993), Chick (1996)) and some of those results are applicable here.

Maximum likelihood methods may be used to estimate the unknown parameters in the function $f(n)$. The log-likelihood for the data is

$$\ell = \sum_{i=1}^{M} x_i \log(h(\mu_i)) + (k_i - x_i) \log(1 - h(\mu_i)) \tag{5}$$

where k_i and x_i are, respectively, the total number of tests and the number of positive responses at dilution level i. Numerical maximization of (5) is required to obtain the mle's $\hat{\theta}$ and $\hat{\phi}$. Variances of these estimates may be obtained by inverting the negative of the matrix of second derivatives of ℓ (see appendix). Cyr et al. (1993) discusses the relative merits of different approaches for generating confidence intervals for θ. If all of the negative control samples test negative, then $\hat{\phi} \approx 1$ and confidence intervals for ϕ based on asymptotic results will not be valid. However, by applying the "rule of three's" (Handley and Lippman-Hand, 1983) a 95% confidence interval for ϕ is given by $[(k_0 - 3)/k_0, 1.0]$ where k_0 is the number of negative control samples. This confidence interval is conservative in the sense that it does not use any of the information from the dilutions where $\mu_i > 0$ for estimating ϕ.

For a given n, the estimated sensitivity (1 - specificity if $n = 0$) is

$$\hat{\alpha} = 1 - \hat{\phi}(1 - \hat{\theta})^n \tag{6}$$

while the value of n that corresponds to a sensitivity of α may be estimated by

$$\hat{n} = \frac{\log(1 - \alpha) - \log(\hat{\phi})}{\log(1 - \hat{\theta})} \tag{7}$$

The asymptotic variances of these estimates are given in the appendix.

2.2   Nonparametric estimation of f

There is no assurance that the parametric model described above will provide a good fit to the curve $f(n)$. As an alternative, isotonic regression methods (Barlow et al., 1972) can be used to provide a nonparametric estimate of an arbitrary non-decreasing function. Typically, however, isotonic regression is used to regress an observed response against an observed covariate. For example, if data were available at each n, then a simple isotonic regression estimate of f could be obtained using the "pool adjacent violators" algorithm. Specifically, let $f(n)$ be the unconstrained estimate of the detection probability when exactly n molecules of DNA are in the sample. That is,

$$f(n) = x_n/k_n \tag{8}$$

Then the isotonic regression estimate of f on n is defined by

$$f^I(n) = \min_{t \geq n} \max_{s \leq n} \sum_{i=s}^{t} k_n f(i) / \sum_{i=s}^{t} k_n \tag{9}$$

Heuristically, this formula sets f^I equal to f for portions of the curve where f is non-decreasing and sets f^I equal to the average of the f where f is decreasing. It can be shown that this algorithm gives the constrained mle of f.

Unfortunately, as we have noted, it is not possible to directly observe x_n and k_n for specific n. However, it is possible to define an iterative estimation scheme using an EM

7

algorithm (Dempster et al, 1978) that combines the usual E-step with an isotonic regression in the M-step. At iteration j the EM algorithm simply replaces equation 8 with

$$f^{(j)}(n) = \frac{\sum_{i=1}^{M} x_i Pr(n \mid x = 1, \mu_i, f^{I(j-1)})}{\sum_{i=1}^{M} x_i Pr(n \mid x = 1, \mu_i, f^{I(j-1)}) + (k_i - x_i) Pr(n \mid x = 0, \mu_i, f^{I(j-1)})} \qquad (10)$$

where $f^{I(j-1)}$ is the current estimate of f^I and

$$Pr(n \mid x, \mu, f) = \frac{f(n)^x (1 - f(n))^{1-x} \mu^n / n!}{\sum_{\nu} f(\nu)^x (1 - f(\nu))^{1-x} \mu^{\nu} / \nu!} \qquad (11)$$

is the conditional probability that the number of molecules in the aliquot is n, given the observed data and current estimate of f^I. The $f^{(j)}$ are then adjusted using the isotonic regression equation (9) to produce a new value of $f^{I(j)}$. The quantity k_n in equation 9, which acts as a weight in the isotonic regression, is replaced by the denominator of equation 10. A convenient starting estimate of f is given by

$$f^0(n) = \frac{\sum_{i=1}^{M} x_i Pr(n \mid \mu_i)}{\sum_{i=1}^{M} Pr(n \mid \mu_i)} \qquad (12)$$

where $Pr(n|\mu_i)$ is the Poisson probability given in (1). Estimation of α for a given n and inverse estimation of n for a given α may be accomplished in a straightforward manner from a plot or table of $f^I(n)$ versus n. The asymptotic properties of this model (in particular standard errors) are unclear, however, since the number of parameters is unlimited.

2.3   Semiparametric estimation of f

As a compromise between the strong assumptions of a parametric model for f and the typically rough nonparametric estimate of f, we consider modeling f as an arbitrary smooth, but monotonic, function of n. While splines are often used to model smooth functions (e.g. Wahba, 1990), modeling a monotone function using standard spline methods would be difficult since the model parameters would need to be constrained to ensure monotonicity. Ramsay (1998), however, developed a simple transformation that guarantees monotonicity

8

of a spline function without constraints on the parameters. Following Ramsay (1998) we define f as

$$\mathrm{logit} f(n) = \beta_0 + \beta_1(\int \exp(\int w(n))) \qquad (13)$$

where $w(n)$ is an arbitrary smooth (spline-based) function (see below for further details). Since the form of w is arbitrary the coefficients in the model for w are unconstrained. Nonetheless, f will be monotonic in n (increasing if $\beta_1 > 0$ and decreasing if $\beta_1 < 0$). In addition, by modeling the logit of f, we also remove any constraints on the β's due to bounds on f.

To fit this function we maximize the penalized log likelihood

$$\ell = \sum_{i=1}^{M} x_i \log(h(\mu_i)) + (k_i - x_i) \log(1 - h(\mu_i)) - \rho \int_0^\infty w^2(n) dn \qquad (14)$$

where ρ controls the smoothness of the fit. For $\rho = 0$, f equals the nonparametric regression estimate given in the previous section (assuming sufficient knots are used in defining w - see below) while $\rho = \infty$ leads to a straight line on the logit scale.

The function $w(n)$ may be defined as a linear combination of an arbitrary set of basis functions,

$$w(n) = \sum_i c_i \phi_i(n).$$

The M-spline basis described in Ramsay (1988) forms a convenient set of basis functions, ϕ, since the integral $\int w(n)$ is then a simple function of the I-spline basis functions, Φ:

$$\int w(n) = \sum_i c_i \Phi_i(n).$$

In deriving the bases, we choose knot positions based on the regions of maximum curvature seen in a plot of the nonparametric estimate of f. A large number of knots are typically included and the smoothing parameter ρ is used to reduce the effective degrees of freedom and avoid overfitting.

A simple cross-validation algorithm may be used to determine the value of ρ. At each dilution level, a single positive result is removed from the data, the model is refit and the

9

misclassification error for the excluded result is computed (i.e. since the excluded result is a positive test, the misclassification error is the probability of a negative test at that dilution level based on equation 2). This procedure is repeated for a negative result at each dilution level. The total misclassification error at a given dilution level is the number of positive results at that dilution level times the misclassification probability for a positive result plus the number of negative results at that dilution level times the misclassification probability for a negative result. The misclassification errors are then summed over the dilution levels. The value of ρ that minimizes the total misclassification error is used in the final model.

Estimation of α for a given n and inverse estimation of n for a given α may be obtained in a straightforward manner from the fitted spline model. Confidence intervals for these estimates may be obtained using the approximate bootstrap confidence interval (ABC) described by Efron and Tibshirani (1993; pg 328). The ABC method has good asymptotic properties and requires fewer model fits than standard bootstrap confidence intervals that rely on resampling.

3   Example

An experiment was designed to estimate the sensitivity and specificity of a new PCR-based test for detecting *M. genitalium*. PCR amplifications were run on 100 μl aliquots of *M. genitalium* at eight dilution levels (mean concentrations of 64, 32, 16, 8, 4, 2, 1 and 0 molecules of DNA per 100 μl). Each aliquot was classified as positive or negative by a microwell detection assay (see Dutro et al., 2002 for further details). The data are summarized in table 1.

¡table 1 here¿

Figure 1 shows the three fitted models for $f(n)$ while table 2 provides estimates of key quantities based on these models. All three curves suggest that the probability of a positive test is substantially less than 1.0 when only a few copies of the target DNA are available. Also, there is no clear detection limit above which tests will always be positive and below

10

which tests will always be negative. The parametric and semiparametric models agree quite well at low copy numbers but then the semiparametric curve rises more quickly and flattens out sooner than the parametric curve, leading to substantial differences in estimates of the minimum detectable copy number for sensitivities above 90%. The nonparametric model is broadly similar to the parametric and semiparametric models but the roughness of the nonparametric curve makes detailed comparisons more difficult.

¡figure 1 here¿

¡table 2 here¿

Figure 2 shows the observed data and the observed sensitivity curves the parametric and semiparametric models. Both agree well with the observed data.

¡figure 2 here¿

4   Discussion

The term sensitivity has been used in a number of ways in the literature on PCR-based tests. However, none of the definitions acknowledge that sensitivity may depend on the number of target cells or molecules in the sample. While this assumption may be reasonable for less sensitive tests that have detection limits in the hundreds or thousands of targets, it seems less reasonable for highly sensitive PCR-based tests. In any event it is an assumption that can be checked with appropriate data. We argue that the sensitivity of a test should be characterized by a curve showing the probability of a positive test as a function of the number of target DNA molecules in the sample. Note that we distinguish the test sensitivity and specificity as defined herein from clinical sensitivity and specificity. The latter would also depend on patient characteristics, the presence of non- target bacteria that might cross-react with the PCR primers, efficiency of lysis of bacteria to free the DNA, specimen preparation and other aspects of the clinical sample collection procedure.

An endpoint dilution experiment may be used to estimate the sensitivity and specificity

11

of a PCR-based test. An important limitation of the endpoint dilution design, however, is that the test sensitivity is confounded with variability in DNA concentration. Nonetheless, by assuming that the number of DNA copies in a particular sample follows a Poisson distribution it is possible to estimate the sensitivity curve, $f(n)$, from the observed counts of positive and negative results at various dilution levels. We describe parametric, nonparametric and semiparametric models for $f(n)$. In each case the resulting curve is constrained to be monotonically increasing.

Generally, use of the parametric model (3) will be computationally fastest and easiest to implement. This model is also easy to interpret in the sense that sensitivity and specificity are each characterized by a single parameter. However, this model assumes a particular functional form for $f(n)$ and data at one dilution level can strongly influence the curve at other, distant dilution levels. In contrast, the effect of the data is much more local in the nonparametric model. However, in our application the nonparametric model gave a relatively "rough" estimate of $f(n)$ and this may generally be the case with the amount of laboratory data typically available. As a compromise between these two approaches, a model based on monotonic smoothing splines can provide a smooth curve for $f(n)$ with minimal assumptions and reduced influence of the data on distant dilution levels. However, the spline-based model is more computationally intensive than the parametric model, especially for computing confidence intervals. Also, comparisons between sensitivity curves from two or more competing tests may be more difficult with the spline-based model. For the data analyzed here, the spline-based model suggested that the parametric model provided a reasonable fit. There is no guarantee that this will be generally true, however.

Thus far, we have assumed that the expected number of DNA targets, μ, is known exactly. In practice, however, there is some uncertainty about the initial concentration of targets and addition uncertainty may be introduced during the process of serial dilution. If the investigator is willing/able to quantify this uncertainty, then it is possible to incorporate it into the analysis. Suppose that the expected number of targets at dilution level i may be

described by a gamma distribution with mean μ_i and squared coefficient of variation equal to $1/c$. Then the conditional distribution of n given μ_i and c is negative binomial,

$$Pr(n \mid \mu_i, c) = \frac{\Gamma(n+c)}{\Gamma(n+1)\Gamma(c)} \left(\frac{\mu_i}{\mu_i + c} \right)^n \left(\frac{c}{\mu_i + c} \right)^c. \tag{15}$$

This distribution may be incorporated into the parametric, nonparametric or semiparametric analyses. If the parametric model (3) holds, the observed probability of a positive test has the simple form

$$h(\mu_i, c) = 1 - \phi \left(1 + \frac{\mu_i}{c}\theta \right)^{-c}. \tag{16}$$

This quantity may be used in the likelihood (5) to obtain estimates of ϕ and θ.

Once sensitivity and specificity have been estimated, the unknown concentration of target DNA in a new sample may be estimated using an LDA experiment. Let λ be the unknown concentration in the new sample and let d_i be the dilution factor for the i'th dilution level (so the concentration at the i'th dilution level is λ/d_i). Then an estimate of λ may be obtained by solving the following equation for λ,

$$\sum_i \frac{x_i \frac{dh\left(\frac{\lambda}{d_i}\right)}{d\lambda}}{h\left(\frac{\lambda}{d_i}\right)} = \sum_i \frac{(k_i - x_i)\frac{dh\left(\frac{\lambda}{d_i}\right)}{d\lambda}}{1 - h\left(\frac{\lambda}{d_i}\right)} \tag{17}$$

If the parametric model (3) is used for $f(n)$ then the estimate of λ is the solution to

$$\sum_i \frac{\hat{\phi} x_i}{d_i \left(\exp(\hat{\theta}\lambda/d_i) - \hat{\phi} \right)} = \sum_i \frac{k_i - x_i}{d_i}. \tag{18}$$

13

References

[1] Barlow RE, Bartholomew DJ, Bremner JM, Brunk HD (1972) Statistical Inference under Order Restrictions. Wiley, Chichester, England.

[2] Chick SE (1996) Bayesian models for limiting dilution assay and group test data. *Biometrics, 52*, 1055 - 1061.

[3] Cyr L, Rust PF, Peters JR, Schmehl MK, Bank HL (1993) Confidence intervals for the relative frequency of responding cells in limiting dilution assays. *Biometrics, 49*, 491 - 498.

[4] Dempster A, Laird N, and Rubin D (1977) Maximum likelihood from incomplete data via the EM algorithm. *J R Statist Soc B, 39*, 1-22.

[5] Dutro SM, Hebb JK, Garin CA, Kenny GE, Hughes JP, Totten PA (2002) Development and performance of a high-throughput PCR assay for *Mycoplasma genitalium.* Submitted to *J. Clinical Microbiology*

[6] Efron B and Tibshirani RJ (1993) An Introduction to the Bootstrap. Chapman and Hall, New York, 436pp.

[7] Handley JA and Lippman-Hand A (1983) If nothing goes wrong, is everything all right? Interpreting zero numerators. *J. Amer. Med. Assoc., 249*, 1743 - 1745.

[8] Ramsay JO (1988) Monotone regression splines in action. *Statistical Science, 3*, 425-441.

[9] Ramsay JO (1998) Estimating smooth monotone functions. *J. R. Statist. Soc. B, 60*, 365-375.

[10] Rodrigo AG, Goracke PC, Rowhanian K, Mullins JI (1997) Quantitation of target molecules from polymerase chain reaction-based limiting dilution assays. *AIDS Research and Human Retroviruses, 13*, 737-742.

14

[11] Shapiro DS (1999) Quality control in nucleic acid amplification methods: Use of elementary probability theory. *J. Clinical MicroBiol., 37*, 848-851.

[12] Sharrock CEM, Kaminski E, Man S (1990) Limiting dilution analysis of human T cells: a useful clinical tool. *Immunology Today, 11*, 281-286.

[13] Sykes PJ, Neoh SH, Brisco MJ, Hughes E, Condon J, Morley AA (1992) Quantitation of targets for PCR by use of limiting dilution. *Biotechniques, 13*, 444-449.

[14] Taswell C (1981) Limiting dilution assays for the determination of immunocompetent cell frequencies. *J. Immunol., 126*, 1614-1619.

[15] Wahba G (1990) Spline Models for Observational Data. CBMS-NSF Regional Conference Series in Applied Mathemathics 59, Society for Industrial and Applied Mathematics, Philadelphia, PA, 169pp.

Appendix

The expected information about $\hat{\phi}$ and $\hat{\theta}$ from model (3) is

$$
\begin{pmatrix}
\sum_i \frac{\mu_i^2 k_i \phi \exp(-\theta\mu_i)}{1-\phi \exp(-\theta\mu_i)} & -\sum_i \frac{\mu_i k_i \exp(-\theta\mu_i)}{1-\phi \exp(-\theta\mu_i)} \\
-\sum_i \frac{\mu_i k_i \exp(-\theta\mu_i)}{1-\phi \exp(-\theta\mu_i)} & \sum_i \frac{k_i \exp(-2\theta\mu_i)}{1-\phi \exp(-\theta\mu_i)}
\end{pmatrix}
\tag{19}
$$

Asymptotic variances are obtained by inverting this matrix. In some applications, it may be reasonable to treat ϕ as fixed at 1.0. In that case,

$$
\mathrm{Var}(\hat{\theta}) = \sum_i \frac{\mu_i^2 k_i \exp(-\theta\mu_i)}{1-\exp(-\theta\mu_i)}
\tag{20}
$$

The delta method may then be used to derive the asymptotic variances of $\hat{\alpha}$ and $\hat{n}$ in equations 6 and 7, respectively,

$$
\mathrm{Var}(\hat{\alpha}) = (1-\hat{\theta})^{2n}\mathrm{Var}(\hat{\phi}) + 2n\hat{\phi}(1-\hat{\theta})^{2n-1}\mathrm{Cov}(\hat{\phi},\hat{\theta}) + n^2\phi^2(1-\hat{\theta})^{2n-2}\mathrm{Var}(\hat{\theta})
\tag{21}
$$

$$
\mathrm{Var}(\hat{n}) = \frac{\mathrm{Var}(\hat{\phi})}{\hat{\phi}^2 \log^2(1-\hat{\theta})} - \frac{2\mathrm{Cov}(\hat{\phi},\hat{\theta}) \log \frac{1-\alpha}{\hat{\phi}}}{\hat{\phi}(1-\hat{\theta})\log^3(1-\hat{\theta})} + \frac{\mathrm{Var}(\hat{\theta})\log^2\frac{1-\alpha}{\hat{\phi}}}{(1-\hat{\theta})^2\log^4(1-\hat{\theta})}
\tag{22}
$$

16

Table 1: Data from an endpoint dilution experiment on *M. genitalium*

Mean number of molecules per 100 ul	Number of replicates	Number positive	Number negative
64	16	16	0
32	16	15	1
16	16	14	2
8	16	15	1
4	16	11	5
2	16	6	10
1	16	5	11
0	22	0	22

Table 2: Estimates of specificity, sensitivity for a given number of DNA copies and detectable number of DNA copies for a given sensitivity (α) for the data in table1.

		sensitivity			Minimum detectable number	
Model	specificity	$n = 1$	$n = 10$	$n = 20$	$\alpha = .50$	$\alpha = .95$
parametric	1.0	0.20	0.89	0.99	3.1	13.4
nonparametric	1.0	0.41	0.91	0.91	3.2	36.6
semiparametric	0.98	0.18	0.90	0.93	2.0	28.6
95% CI for	(0.96 -	(0.05 -	(0.79 -	(0.82 -	(1.3 -	(6.2 -
semiparametric	0.99)	0.42)	0.97)	0.99)	2.9)	50.6)
estimates based						
on ABC method						

18