Published in final edited form as:

Acad Radiol. 2013 July ; 20(7): 838–846. doi:10.1016/j.acra.2013.04.001.

Multivariate Normally Distributed Biomarkers Subject to Limits of Detection and Receiver Operating Characteristic Curve Inference

Neil J. Perkins, PhD, Enrique F. Schisterman, PhD, and Albert Vexler, PhD

Division of Epidemiology, Statistics and Prevention Research, Eunice Kennedy Shriver National Institute of Child Health and Human Development, National Institutes of Health, DHHS, 6100 Executive Boulevard, Bethesda, MD 20852 (N.J.P., E.F.S., A.V.); and Department of Biostatistics, State University of New York Buffalo, Buffalo, NY (A.V.)

Abstract

Rationale and Objectives—Biomarkers are of ever-increasing importance to clinical practice and epidemiologic research. Multiple biomarkers are often measured per patient. Measurement of true biomarker levels is limited by laboratory precision, specifically measuring relatively low, or high, biomarker levels resulting in undetectable levels below, or above, a limit of detection (LOD). Ignoring these missing observations or replacing them with a constant are methods commonly used although they have been shown to lead to biased estimates of several parameters of interest, including the area under the receiver operating characteristic (ROC) curve and regression coefficients.

Materials and Methods—We developed asymptotically consistent, efficient estimators, via maximum likelihood techniques, for the mean vector and covariance matrix of multivariate normally distributed biomarkers affected by LOD. We also developed an approximation for the Fisher information and covariance matrix for our maximum likelihood estimations (MLEs). We apply these results to an ROC curve setting, generating an MLE for the area under the curve for the best linear combination of multiple biomarkers and accompanying confidence interval.

Results—Point and confidence interval estimates are scrutinized by simulation study, with bias and root mean square error and coverage probability, respectively, displaying behavior consistent with MLEs. An example using three polychlorinated biphenyls to classify women with and without endometriosis illustrates how the underlying distribution of multiple biomarkers with LOD can be assessed and display increased discriminatory ability over naïve methods.

Conclusions—Properly addressing LODs can lead to optimal biomarker combinations with increased discriminatory ability that may have been ignored because of measurement obstacles.

Keywords

Area under the curve; left censoring; limit of detection; maximum likelihood; ROC

The use of biomarkers, including medical images, is becoming increasingly common in the diagnosis and prognosis of individuals with a given disease in clinical settings and epidemiologic research. Combining various sources of biologic information quantitatively, through diagnostic indices for example, allows a concrete way to collapse blood, imaging, and other biomarker information. The receiver operating characteristic (ROC) curve is a tool frequently used to assess the usefulness of biomarkers as diagnostic tools (1–4). Some biomarkers, often emerging ones, are difficult to measure and/or are measured with less than optimally refined laboratory methods, making it difficult to assess any potential clinical usefulness. Specifically, limits of quantifiable biomarker levels are introduced by the sensitivity of the measurement instrument or process, which often leads to censored measurements at either the lower or upper end of the biomarker spectrum. The points at which this censoring occurs are called the limit of detection (LOD) and point of saturation, respectively (5). However, scientists remain interested in the biomarkers' potential relations, causations, and discriminatory ability for adverse outcomes and disease despite this measurement-induced error of censoring. For instance, polychlorinated biphenyl (PCB) levels are a group of compounds that have been linked with several adverse outcomes but their correlated nature and measurements being affected by LODs have hindered their study.

Consider p biomarkers that are of interest in a particular population, where each individual has true values for each, denoted by $\vec{W} = (W_1, \dots, W_p)^T$. Now suppose laboratory measurements, $\vec{Z}$, for each of those same biomarkers are affected by LODs $\vec{d} = (d_1, \dots, d_p)^T$. One way to write this component-wise transformation of $\vec{W}$ is

$$Z_l = \begin{cases} W_l & ; \quad W_l \geq d_l \\ \cdot & ; \quad W_l < d_l \end{cases},$$

where for a fixed d_l , $l = 1, \dots, p$, the l th biomarker level is either quantified above d_l or censored below d_l . Clearly, estimating the underlying distribution or parameters and testing hypothesis intended for $\vec{W}$ but using $\vec{Z}$, a mix of continuous and ordinal data, requires special consideration.

Complete case analysis, omitting the censored observations from $\vec{Z}$, surely leads to bias and loss of efficiency in any conceivable situation. More frequently, a replacement value like $d/2$ is used for the censored observations in the data. This technique provides a less biased assessment than omission (6). Although unbiased in special cases, replacement with constants has been shown to lead to a biased estimation of mean and standard deviation parameters, regression coefficients, odds ratio, and area under the ROC curve (AUC) for a single biomarker (7–11).

If we can assume distributions of each of the p biomarkers are normal or γ , we could apply univariate parametric results proposed by Gupta (12) and Harter and Moore (13), respectively. These methods estimate the parameters governing the underlying distributions of each W_l through maximizing likelihoods given data measured with an LOD (12–14). For a single biomarker, these developments are sound but they assume the p biomarkers to be independent, which is often inappropriate or at least inefficient for biomarkers of a common

outcome and/or biomarkers that are measured simultaneously. Lyles et al (8) considered exactly two correlated biomarkers affected by LODs where the true levels follow a bivariate normal distribution. They constructed and maximize the appropriate likelihood function to estimate the means, variances, and correlation for exactly $p = 2$ biomarkers.

As multiplex assays and three-dimensional imaging become increasingly popular, the availability of multiple ($p \geq 2$) bio-markers has become the rule rather than the exception and demands methods to be able to use such data. In Materials and Methods, we consider the general case of $p \geq 2$ biomarkers following a multivariate normal distribution, $\vec{W} \sim N_p(\vec{\mu}, \vec{\Sigma})$, where all of the biomarkers are potentially measured with LODs and $\vec{Z}_Y$ is observed rather than $\vec{W}$. The likelihood function for $p \geq 2$ developed here, which simplifies to the likelihood function in Lyles et al when $p = 2$, can be maximized to yield maximum likelihood estimations (MLEs) $\hat{\vec{\mu}}$ and $\hat{\vec{\Sigma}}$ (8,15). To accompany these point estimates, we use Taylor expansion to derive the Fisher information matrix in this setting.

These MLEs can be used for a variety of inference and secondary estimators. Also in the next section, we focus on the ROC curve, a statistical tool commonly used to evaluate the use of biomarkers in discriminating between diseased and healthy populations. In this multiple biomarker setting, we use data measured with LODs to construct an index that is the best linear combination (BLC) of the p biomarkers, where “best” means the combination that maximizes the AUC (16). Assuming multivariate normality, we show how to first

obtain the $\hat{\vec{\mu}}_X$ and $\hat{\vec{\Sigma}}_X$ and the $\hat{\vec{\mu}}_Y$ and $\hat{\vec{\Sigma}}_Y$ for diseased and healthy populations, respectively. Subsequently, these MLEs can be used to estimate the BLC, an ROC curve, and accompanying AUC with approximate $1 - \alpha$ level confidence intervals (CIs) for the AUC.

Point and CI estimators are evaluated via simulation in the Results. In Example, we use levels of $p = 3$ PCBs, measured with LODs, to classify women with and without endometriosis to illustrate the estimation of means and covariance, the BLC of the three PCBs, and the resulting ROC curve with AUC. We end with a brief discussion of issues surrounding estimation and ROC analysis based on multiple biomarkers affected by LODs.

MATERIALS AND METHODS

Left-Censored Multivariate Normal Likelihood

Suppose that biomarker levels $\vec{W}$ follow a multivariate distribution $g(\vec{W}; \vec{\theta})$ but measurements $\vec{Z}$ are actually observed with fixed LODs $\vec{d} = (d_1, \dots, d_p)^T$. Clearly, in most cases a likelihood function for parameters $\vec{\theta}$ developed for the random sample $\vec{W}$ would be inappropriate for sample $\vec{Z}$. This likelihood would not account for the censored values, even when censored values are replaced by a constant. A likelihood function can be formed that does account for censoring and maximizing the appropriate likelihood results in MLEs that are consistent, asymptotically unbiased. To build this likelihood, let $\vec{Q}(\vec{Z})$ be a p -dimensional vector of indicator functions, $Q_l = I(Z_l \geq d_l)$, that take the value 1 when the measured bio-marker is greater than or equal to the LOD and 0 when it is less than the LOD. This creates, 2^p different $\vec{Q}(\vec{Z})$ that are possible and can be observed if each biomarker is measured with LOD. The extremes $\vec{Q}(\vec{Z}) = (1, 1, 1, \dots, 1)$, indicating no censoring, and $\vec{Q}(\vec{Z}) = (0, 0, 0, \dots, 0)$, indicating all biomarkers are censored, are also possible.

$(\vec{Z}) = (0, 0, 0, \dots, 0)$, indicating that all of that individual's measurements are censored, set the edges of a variety of mixed observed and censored vectors. Individuals with the same pattern of censoring, $\mathbf{Q}(\vec{Z})$, have the same likelihood form. Thus, we can use these different patterns of censoring to construct a likelihood for our entire sample to maximize. The likelihood for the bivariate normal distribution in Lyles et al (8) had $2^2 = 4$ parts for $p = 2$.

Here we are considering the case where the biomarkers, W are multivariate normally distributed with mean vector μ and covariance matrix Σ . Suppose that we only observe the measured $\vec{Z}$. Let us start creating our likelihood by first looking at the $\mathbf{Q}(\vec{Z}) = (1, 1, 1, \dots, 1)$, where the likelihood function is the standard result $L(\mu, \vec{\Sigma}; \vec{Z}) = L_1(\mu, \vec{\Sigma}; w_1, w_2, \dots, w_p) = g(w; \mu, \vec{\Sigma})$. The other extreme, where all biomarkers are censored, $\mathbf{Q}(\vec{Z}) = (0, 0, 0, \dots, 0)$, corresponds to a likelihood, $L(\mu, \vec{\Sigma}; \vec{Z}) = L_{2^p}(\mu, \vec{\Sigma}; w_1 < d_1, w_2 < d_2, \dots, w_p < d_p)$, that is the simple probability that all biomarkers are less than the LODs. While the details are presented in the Appendix, the probability that all of an individual's biomarker levels are below the LODs for a given μ and Σ will be denoted by $A(\vec{d}; \mu, \vec{\Sigma})$. For the likelihoods of all the cases of $\mathbf{Q}(\vec{Z})$ in between, $L_k(\mu, \vec{\Sigma}; \vec{Z})$, $k = 2, 3, \dots, 2^p - 1$; that is, those $\vec{Z}$ composed of both measured and censored levels, we extend ideas from explicit results for $p = 2$ to the general p case here (8,15). In that work, the bivariate normal case with $\mathbf{Q}(\vec{Z}) = (1, 0)$ or $(0, 1)$ was shown to lead to a likelihood that is the product of a marginal distribution for the biomarker that is observed and a conditional distribution for the biomarker that is censored. We will rely on the same rationale here in a multivariate framework.

All mixed vectors can be reordered and partitioned to have the same general form such that $\vec{Z} = (\vec{Z}^{(1)}, \vec{Z}^{(2)})$, where the leading set is the observed biomarkers and the latter is the censored. The resultant multivariate normal density is then factored into a marginal density, $g_{(1)}(\vec{W}^{(1)}; \mu^{(1)}, \Sigma_{11})$, and a conditional density, $g_{(2|1)}(\vec{W}^{(2)}; \mu^{(2)}, \Sigma_Q)$, with

$\vec{\mu}^{(Q)} = \vec{\mu}^{(2)} + \sum_{21} \sum_{11}^{-1} (\vec{W}^{(1)} - \vec{\mu}^{(1)})$ and $\Sigma_Q = \Sigma_{22} - \sum_{21} \sum_{11}^{-1} \sum_{12}$. Since the observations in $\vec{Z}^{(2)}$ are censored below the LODs, their contribution to the likelihood is simply the probability that those biomarkers would be below those LODs. Each of the $k = 2, 3, \dots, 2^p - 1$ likelihood contributions has the same form

$$L_k \left(\vec{u}, \Sigma; \vec{Z}^{(1)}, \vec{Z}^{(2)} \right) \propto g_{(1)} \left(\vec{w}^{(1)}; \vec{\mu}^{(1)}, \Sigma_{11} \right) \times A_{(2|1)} \left(\vec{d}^{(2)}; \vec{\mu}^{(Q)}, \Sigma_Q \right).$$

So for a random sample $\mathbf{Z}$ of independent identically distributed p -dimensional vectors $\vec{Z}_j$, $j = 1, \dots, n$, following a multivariate normal distribution, the likelihood has the form

Maximizing Equation 1 with respect to each element of μ and Σ results in MLEs $\hat{\mu}$ and $\hat{\Sigma}$. This must be performed numerically due to the complexity of the likelihood for which a closed-form solution does not exist for even the univariate case. The details of this development are found in the Appendix.

$$\begin{aligned}
L(\vec{\mu}, \Sigma; \mathbf{Z}) &= \prod_{j=1}^{n_Y} L(\vec{\mu}, \Sigma; \vec{\mathbf{Z}}_j) \\
&\propto \prod_{\vec{\mathbf{Q}}(\vec{\mathbf{Z}}_j) = (1, 1, \dots, 1)^T} g(\vec{\mathbf{w}}_j; \vec{\mu}, \Sigma) \times \dots \\
&\times \prod_{\vec{\mathbf{Q}}(\vec{\mathbf{Z}}_j) = (1, 0, \dots, 1, 1, 0)^T} g_{(1)}(\vec{\mathbf{w}}_j^{(1)}; \vec{\mu}^{(1)}, \Sigma_{11}) \times \Lambda_{(2|1)}(\vec{\mathbf{d}}^{(2)}; \vec{\mu}^{(Q)}, \Sigma_Q) \times \dots \quad (1) \\
&\times \prod_{\vec{\mathbf{Q}}(\vec{\mathbf{Z}}_j) = (0, 0, \dots, 0, 0)^T} \Lambda(\vec{\mathbf{d}}; \vec{\mu}, \Sigma)
\end{aligned}$$

To accompany these point estimates, we used Taylor expansion in the Appendix to generate covariance matrices that will allow us to create CIs for the AUC. Of note is that it is well known that $\hat{\vec{\mu}}$ and $\hat{\Sigma}$ are independent in the setting without LODs. We have lost that property thanks to the introduction of the conditional probability into the likelihood when some of the data are below the LOD.

ROC Curve and the AUC

Suppose the multiple biomarkers measured with LODs are intended to be used as diagnostic indices of disease status for a particular outcome. The ROC curve is used to assess such diagnostic ability by plotting sensitivity [the probability of a positive test result given the individual is truly diseased, denoted by $q(c)$] by $1 - \text{specificity}$ [the probability of a negative test result given the individual is truly healthy, denoted by $p(c)$] across all possible cut points c . Often a positive or negative test is indicated by having a biomarker level above or below, respectively, a given c but can be similarly used when the opposite is true. The ROC curve can be used to assess discriminatory ability over all c , over a specific range of $q(c)$ or $p(c)$, and to identify cut points that result in a specific discriminatory ability (3,4). The AUC, or $P(X > Y)$, is the most commonly used summary measure, ranging from 0.5 to 1, with larger values indicating a greater ability to discriminate (3,4). Consideration has already been given to its estimation based on a single biomarker affected by LOD (11). Su and Liu have shown that a linear combination can achieve discrimination greater than using a single biomarker alone (16).

Let p biomarkers for individuals with and without a particular disease follow multivariate normal distributions, $\vec{\mathbf{X}} \sim N_p(\vec{\mu}_X, \Sigma_X)$ and $\vec{\mathbf{Y}} \sim N_p(\vec{\mu}_Y, \Sigma_Y)$, respectively. Then, indices formed

from the linear combinations of those biomarkers, $U = \vec{\beta}^T \vec{\mathbf{X}} = \sum_{j=1}^p \beta_j X_j$ and

$V = \vec{\beta}^T \vec{\mathbf{Y}} = \sum_{j=1}^p \beta_j Y_j$, are univariate normally distributed, $U \sim N_1(\vec{\beta}^T \vec{\mu}_X, \vec{\beta}^T \Sigma_X \vec{\beta})$ and $V \sim N_1(\vec{\beta}^T \vec{\mu}_Y, \vec{\beta}^T \Sigma_Y \vec{\beta})$, respectively. Consequentially, an ROC curve can be formed from U and V that depends on the choice of $\vec{\beta}$. The choice of $\vec{\beta}$ can be equally weighted or an alternative designated by expert knowledge. In this setting, Su and Liu developed the BLC of multiple biomarkers

$$\vec{\beta}_0^T \propto (\vec{\mu}_X - \vec{\mu}_Y)^T (\Sigma_X + \Sigma_Y)^{-1}, \quad (2)$$

leading to

$$= \Phi \left(\frac{AUC_0 = P(U > V)}{\sqrt{(\hat{\mu}_X - \hat{\mu}_Y)^T (\Sigma_X + \Sigma_Y)^{-1} (\hat{\mu}_X - \hat{\mu}_Y)}} \right), \quad (3)$$

which is “best” with respect to maximizing AUC over all real β^T (16).

Using samples measured with LOD, $\mathbf{Z}_X$, and $\mathbf{Z}_Y$, we can maximize the likelihood in

Equation 1 for both groups, generating MLEs $\hat{\mu}_X$ and $\hat{\Sigma}_X$ and $\hat{\mu}_Y$ and $\hat{\Sigma}_Y$. Placing these MLEs in Equations 2 and 3 results in MLEs $\hat{\beta}_0^T$ and $\hat{AUC}_0$. Without loss of generality, we assume that both diseased and nondiseased biomarker levels are affected by the censoring at the same level, $d_{xI} = d_{yI} = d_I$. The development of the CI for AUC_0 is shown in the Appendix (17).

RESULTS

The MLEs developed here are asymptotically unbiased. The small sample performance of these estimators was evaluated via simulation in the R programming language and through recorded bias and root mean square error (RMSE) over $B = 2000$ iterations. Coverage probabilities and width of 95% CIs for AUC_0 were recorded. Increasing sample sizes, $n_x = n_y = 50, 100, 200$, were used to demonstrate estimator performance.

Nondiseased values were generated from a $p = 3$ multivariate normal distribution with mean zero and unit diagonal covariance matrices with uniform correlation between bio-markers of $\rho = 0.5$. Additional simulations were conducted with a higher uniform correlation between biomarkers of $\rho = 0.8$. Diseased values were similarly obtained but with a mean μ_X^T

corresponding to $AUC_0 = 0.6, 0.7, 0.8$. For each simulated group of samples, $\hat{\mu}_X, \hat{\Sigma}_X, \hat{\mu}_Y$, and $\hat{\Sigma}_Y$ were calculated, leading to $\hat{\beta}_0^T, \hat{AUC}_0$, and 95% CI for AUC_0 . Due to the complexity of the elements of $Cov(\hat{\mu}, \hat{\Sigma})$ in Equation 2, occasional numeric integration was necessary within some of the elements or portions of the elements.

We calculated these estimators while applying various levels of censoring, where d values corresponded to the quantiles of the marginal distributions of the nondiseased population. First “even” censoring was invoked by applying a similar d to all resulting biomarkers. Levels of even censoring were set at the 10th, 25th, and 50th quantiles across all biomarkers. To assess uneven censoring, d values were chosen at the following quantile sets: (10th, 10th, 50th), (10th, 25th, 50th), and (10th, 50th, 75th).

As expected, the bias and RMSE of $\hat{AUC}_0$ decrease as sample sizes increase. The relative bias, $(\hat{AUC}_0 - AUC_0)/AUC_0$, ranged from 0.010 to 0.152, 0.003 to 0.088, and -0.0001 to 0.050 for $n_x = n_y = 50, 100$, and 200, respectively, including even the highest levels of censoring. Ranges of the RMSE were 0.0433 to 0.1443, 0.0305 to 0.0892, and 0.0213 to

0.0546 for $n_x = n_y = 50, 100$, and 200 , respectively. The majority of the relative biases and RMSEs tended to the lower end of these ranges in all sample sizes. Figure 1 depicts the relative bias and RMSE for the $\rho = 0.5$ and $AUC_0 = 0.6$ (triangles) and 0.8 (squares) cases across all scenarios of censoring and as sample size increases $n_x = n_y = 50, 100$, and 200 (corresponding to increasing point size). The first column of points of each plot is for the MLE with no censoring and serves as a reference for subsequent MLEs based on data with censoring. Clearly within each column both relative bias and RMSE decrease as function of AUC_0 (same-size squares lower than triangles) and sample size (increasing size of same shape). Also evident is that relative bias and RMSE increase as censoring increases, generally from left to right. Figure 1 shows that relative bias and RMSE are only slightly increased when comparing columns 2, 3, 5, and 6, which have variations of 10% and 25% censoring, to column 1 based on the full data. Columns 4 and 7 show that while relative bias and RMSE are markedly increased from column 1, estimation is still potentially useful, accounting for 50% and 75% censored biomarker data. This is especially true for $n_x = n_y = 100$ and 200 near the middle and bottom, respectively, of both plots. While not displayed here, results for $AUC_0 = 0.7$ generally lie in between points in Figure 1, and results for $\rho = 0.8$ show precisely the same relationships with a small increase in magnitude of relative bias and RMSE.

The coverage probability of 95% CI's were nominal or near nominal coverage for all AUC_0 averaging 0.940, 0.945, and 0.945 for $n_x = n_y = 50, 100$, and 200 , respectively. No discernible pattern exists regarding $\rho = 0.5$ and 0.8 or with regard to censoring patterns due to the LOD.

EXAMPLE

Endometriosis is a gynecologic disease exclusive to menstruating species such as humans and primates and occurs primarily in reproductive-age women. Data from experimental animal, primate, and human studies suggests an association between dioxin and PCBs and endometriosis (18). Here, we chose estrogenic congeners PCBs 188, 153, and 180 with LODs $d_{188} = 0.020$, $d_{153} = 0.200$, and $d_{180} = 0.034$ to illustrate our method, by exploring these environmental toxicants as potential indicators of endometriosis using biomarker levels on 28 cases and 51 controls. These LODs of PCBs 188, 153, and 180 resulted in censoring of 32%, 64%, and 0% of cases and 36%, 74%, and 11% of controls, respectively.

Considering each biomarker alone, the appropriateness of univariate normal distributions was assessed via the use of QQ plots (Fig 2) on the log-transformed biomarkers, similar to Lyles et al (8). Diagonal points corresponding to the levels above the LODs shows the three biomarkers fit log normal distributions well based on the quantifiable samples. The horizontal points are essentially quantile placeholders for the observations below the LOD. Assuming that the data we do not see follow the data we do see, univariate log normal likelihoods led to the MLEs $\hat{AUC}_{188} = 0.518$, $\hat{AUC}_{153} = 0.511$, and $\hat{AUC}_{180} = 0.630$ (11). PCB 180 has a markedly higher estimated ability to discriminate and might be the choice if a single marker had to be chosen.

Individuals are, however, exposed to these biomarkers as mixtures and are closely linked biologically, leading to them being highly correlated. Thus, using multiple biomarkers jointly makes sense here. Current practice is simply to sum the congeners, producing a composite exposure but one that is overly simplified (19). We will use developments from Materials and Methods to construct a more effective combination for the outcome endometriosis.

Applying our developments from Materials and Methods to the vector of log-transformed PCBs 188, 153, and 180, we maximize Equation 1 after grouping the $\mathbf{Z}_{X_i}$ and $\mathbf{Z}_{Y_j}$ by pattern of censoring, $\mathbf{Q}(\mathbf{Z})$, to generate MLEs

$$\hat{\boldsymbol{\mu}}_X = (-3.68, -1.81, -2.23) \quad \hat{\boldsymbol{\mu}}_Y = (-3.72, -1.89, -2.51) \\ \hat{\Sigma}_X = \begin{bmatrix} 0.30 & 0.16 & 0.14 \\ 0.16 & 0.34 & 0.26 \\ 0.14 & 0.26 & 0.23 \end{bmatrix} \quad \hat{\Sigma}_Y = \begin{bmatrix} 0.22 & 0.15 & 0.22 \\ 0.15 & 0.15 & 0.25 \\ 0.22 & 0.25 & 0.45 \end{bmatrix}.$$

While these mean and covariance estimates are informative on their own, we substitute them into Equation 3 and get the BLC $\hat{\boldsymbol{\beta}}_0^T = (0.164, -1.496, 1.598)$. We could now apply these coefficients and generate a parametric ROC curve for the BLC, which we compare to a similar curve for PCB 180 in Figure 3. It is obvious in Figure 3 that the estimated BLC ROC curve (solid) is completely and substantially higher than that of estimated PCB 180 (dashed) and implies a potential increase in discrimination. We can quantify this using the $A\hat{U}C_0$ by simply using Equation 4 to see what might be achieved using all three biomarkers jointly for separating women with and without endometriosis. Combining PCBs 188, 153, and 180 in this manner yields $A\hat{U}C_0 = 0.731$, an increase in estimated discriminatory ability from any single biomarker. Using Equations A2 and A5, we can also generate a 95% CI of 0.677 to 0.779. While these estimates show a potential benefit of using three biomarkers over one, ROC analysis is a process in which superiority of one biomarker over another for practical purposes needs to be validated with independent data before a recommendation can be made.

If we use conventional naïve replacement methods and substitute the d for values below the LOD, we can also obtain a naïve BLC. The linear combination will certainly not be “best” in the sense of maximizing AUC because the assumption of multivariate normality is clearly violated. However, we generated such an estimate, $A\hat{U}C_{0R} = 0.652$ for the sake of comparison, which is significantly less than our $A\hat{U}C_0 = 0.731$.

DISCUSSION

As exposure assessment continues to stray from a single cause to mixtures of chemicals, images, and biomarkers, multiple bio-markers are increasingly being assessed and often measured simultaneously. However, whether measuring single or multiple biomarkers at once, emerging or established limitations of laboratory sensitivity resulting in operational limits, both upper and lower, are nothing new. The methods developed here provide a

solution to the very common real-world data problem of censored values below the LOD, allowing for improved assessment of the biomarkers' joint distribution and countless secondary analyses like regression, Bayesian inference, and, as we have shown, the ROC curve. While most imaging data are not left censored, continuous imaging variables may be part of a combination along with other biomarkers that are as censored.

Multivariate normality is a widely used distribution and is often a reasonable assumption. Here, as in univariate literature, we rely on the assumption that the distribution of the data that we see above the LOD appropriately characterizes the data below that we do not see. This assumption allows for the application of maximum likelihood methods, whose estimator's asymptotic properties, consistency and efficiency, are well known. We used a likelihood function that properly accounted for missing data due to LODs for $P = 2$ biomarkers to generate MLEs for the mean vector and covariance matrix and further used it to develop the Fisher information for such parameters. The developments and simulation study for $A\hat{U}C_0$ demonstrated how these MLEs could be used to generate secondary estimators, which are also consistent and use the Fisher information to create accompanying CIs that achieve nominal coverage probability for a sample size as small as 50 with up to 50% missing values due to LODs. In our simulations, setting the true variances to be equal displayed the performance corresponding to the common setting where a mean shift occurs based on differing disease status. The behavior of $A\hat{U}C_0$, bias, and variability, will differ based on each biomarker scenario and unequal variances, where our simulation might be conservative.

As with all parametric estimation, the appropriateness of distributional assumptions is critical to the reliability of the estimators generated. The literature has shown MLEs of $A\hat{U}C$ for the univariate case are robust to minor deviations from the normal assumption and unfavorable results when assumptions are grossly violated (11,20). Similar caution and diligence should be taken when using the methods developed here for $\hat{\mu}_x$, $\hat{\mu}_y$, $\hat{\Sigma}_x$, $\hat{\Sigma}_y$, and, subsequently, $A\hat{U}C_0$. In our example, we looked at marginal QQ plots to assess the assumption of normality of each PCB and, subsequently, assumed multivariate normality. Regarding the latter assumption, there are cases where multivariate normality fails to follow marginally normally distributed variables, but, practically speaking, these cases are rare and obvious. Care should always be taken to plot and investigate bivariate and multivariate data before any analyses. However, when inference that relies on the distribution of the data below the LOD is of interest, assumptions regarding the behavior of the biomarker values must be made. Using a replacement value makes an assumption that is almost assuredly false by creating a point mass at the chosen value a that grossly misrepresents the biomarkers distribution below the LOD.

The likelihood function developed here can also be used in other applications where biomarkers are measured with LODs and normality can be assumed. Specifically, the ratio of likelihoods might be used to create an indices resulting in a proper ROC curve, without a hook such as in Figure 3, which has been shown to be important in the assessment of medical imaging (21,22). Additionally, our accompanying covariance developments lead to nominal coverage of CIs. A simpler, alternative might be to use the Hessian generated

during the maximization of the likelihood in lieu of our more complex findings. However, a thorough comparison of these two techniques should be conducted in future work before recommending a particular method. Future developments are also necessary to determine if the increased performance of combined biomarkers over that of a single biomarker, similar to that in our example, is statistically significant.

ROC curves are of ever-increasing importance in reader studies of radiologic devices and clinical diagnostics in general (23,24). These methods for estimating $\bar{\mu}_X$, $\bar{\mu}_Y$, $\bar{\Sigma}_X$, $\bar{\Sigma}_Y$, and an ROC curve with AUC_0 while properly accounting for missing data censored below an LOD will provide a consistent view of the potential discriminatory ability of the BLC of a set of biomarkers. Moreover, the binormal ROC curve of the BLC indices means that we can combine methods here for multiple biomarkers with LODs with other important ROC analysis techniques (3,24,25). Assessing these biomarkers and diagnostic data properly will lead to better assessments of exposure and outcome relations. The hope is that we may find biomarkers that were once deemed useless displaying promise that warrants additional resources to improve the measurement process and subsequently lead to improved diagnostic care.

Acknowledgments

This research was supported by the Intramural Research Program of the NIH, Eunice Kennedy Shriver National Institute of Child Health and Human Development.

References

1. Metz CE. Basic principles of ROC analysis. *Semin Nucl Med.* 1978; 8:283–298. [PubMed: 112681]
2. Wagner RF, Beiden SV, Metz CE. Continuous versus categorical data for ROC analysis: some quantitative considerations. *Acad Radiol.* 2001; 8:328–334. [PubMed: 11293781]
3. Pepe, MS. The statistical evaluation of medical tests for classification and prediction. New York: Oxford University Press; 2003.
4. Zhou, XH.; Obuchowski, NA.; McClish, DK. Statistical methods in diagnostic medicine. New York: Wiley & Sons Interscience; 2002.
5. Browne RW, Whitcomb BW. Procedures for determination of detection limits: application to high-performance liquid chromatography analysis of fat-soluble vitamins in human serum. *Epidemiology.* 2010; 21(Suppl):S4–S9. [PubMed: 20220523]
6. Hornung RW, Reed LD. Estimation of average concentration in the presence of nondetectable values. *Appl Occup Environ Hyg.* 1990; 5:46–51.
7. Haas CH, Scheff PA. Estimation of averages in truncated samples. *Environ Sci Technol.* 1990; 24:912–919.
8. Lyles RH, Williams JK, Chuachoowong R. Correlating two viral load assays with known detection limits. *Biometrics.* 2001; 57:1238–1244. [PubMed: 11764265]
9. Schisterman EF, Little RJ. Opening the black box of biomarker measurement error. *Epidemiology.* 2010; 21(Supplement):S1–S3. [PubMed: 20539119]
10. Singh A, Nocerino J. Robust estimation of mean and variance using environmental data sets with below detection limit observations. *Chemometr Intelligent Lab Syst.* 2002; 60:69–86.
11. Perkins NJ, Schisterman EF, Vexler A. Receiver operating characteristic curve inference from a sample with a limit of detection. *Am J Epidemiol.* 2007; 165:325–333. [PubMed: 17110640]
12. Gupta AK. Estimation of the mean and standard deviation of a normal population from a censored sample. *Biometrika.* 1952; 39:260–273.

13. Harter HL, Moore AH. Asymptotic variances and covariances of maximum-likelihood estimators, from censored samples, of the parameters of Weibull and gamma populations. *Ann Math Stat.* 1967; 38:557–570.
14. Vexler A, Liu A, Eliseeva E, et al. Maximum likelihood ratio tests for comparing the discriminatory ability of biomarkers subject to limit of detection. *Biometrics.* 2008; 64:895–903. [PubMed: 18047527]
15. Perkins NJ, Schisterman EF, Vexler A. ROC curve inference for best linear combination of two biomarkers subject to limits of detection. *Biometr J.* 2011; 53:464–476.
16. Su JQ, Liu JS. Linear combinations of multiple diagnostic markers. *J Am Stat Assoc.* 1993; 88:1350–1355.
17. Reiser B, Faraggi D. Confidence intervals for the generalized ROC criterion. *Biometrics.* 1997; 53:192–202.
18. Louis G, Weiner JM, Whitcomb BW, et al. Environmental PCB exposure and risk of endometriosis. *Hum Reprod.* 2005; 20:279–285. [PubMed: 15513976]
19. Cooke, PS.; Sato, T.; Buchanan, DL. Disruption of steroid hormone signaling by PCBs. In: Robertson, LW.; Hansen, LG., editors. *PCBs: recent advances in environmental toxicology and health effects.* Lexington: The University Press of Kentucky; 2001. p. 257-263.
20. Molodianovitch K, Faraggi D, Reiser B. Comparing the areas under two correlated ROC curves: parametric and non-parametric approaches. *Biometr J.* 2006; 48:745–757.
21. Pesce LL, Metz CE, Berbaum KS. On the convexity of ROC curves estimated from radiological test results. *Acad Radiol.* 2010; 17:960–968. [PubMed: 20599155]
22. Pesce LL, Horsch K, Drukker K, et al. Semiparametric estimation of the relationship between ROC operating points and the test-result scale: application to the proper binormal model. *Acad Radiol.* 2011; 18:1537–1548. [PubMed: 22055797]
23. Samuelson F, Gallas BD, Myers KJ, et al. The importance of ROC data. *Acad Radiol.* 2011; 18:257–258. author reply 259–261. [PubMed: 21232688]
24. Alemayehu D, Zou KH. Applications of ROC analysis in medical research: recent developments and future directions. *Acad Radiol.* 2012; 19:1457–1464. [PubMed: 23122565]
25. Hillis SL, Metz CE. An analytic expression for the binormal partial area under the ROC curve. *Acad Radiol.* 2012; 19:1491–1498. [PubMed: 23122568]

APPENDIX

Multivariate Normal Likelihood with Left Censoring

Consider the case where p biomarker levels are multivariate normally distributed,

$$g(\vec{W}; \vec{\mu}, \Sigma) = \frac{1}{(2\pi)^{p/2} |\Sigma|^{1/2}} \exp \left\{ -\frac{1}{2} (\vec{W} - \vec{\mu})^T \Sigma^{-1} (\vec{W} - \vec{\mu}) \right\},$$

where the covariance matrix Σ consists of variance, $\Sigma_{ll} = \sigma_{ll}$, and covariance terms, $\Sigma_{ll'} = \sigma_{ll'}$, $l \neq l'$. Again, suppose that biomarker measurements $\mathbf{Z}$ are observed. The first part necessary for creating a likelihood function for a sample matrix $\mathbf{Z}$ with which we can estimate $\vec{\mu}$ and Σ is for those $\mathbf{Z}$ with $\mathbf{Q}(\mathbf{Z}) = (1, 1, 1, \dots, 1)$, where the likelihood function is the standard result $L(\vec{\mu}, \vec{\Sigma}; \mathbf{Z}) = L_1(\vec{\mu}, \vec{\Sigma}; w_1, w_2, \dots, w_p) = g(w; \vec{\mu}, \vec{\Sigma})$. Second, we skip to the other extreme for those $\mathbf{Z}$ with $\mathbf{Q}(\mathbf{Z}) = (0, 0, 0, \dots, 0)$, where the likelihood function corresponds to the probability

$$\begin{aligned}
 L(\vec{\mu}, \Sigma; \vec{Z}) &= L_{2^p}(\vec{\mu}, \Sigma; w_1 < d_1, w_2 < d_2, \dots, w_p < d_p) \\
 &= \int_{-\infty}^{d_1} \int_{-\infty}^{d_2} \cdots \int_{-\infty}^{d_p} g(\vec{w}; \vec{\mu}, \Sigma) dw_p \cdots dw_2 dw_1 \\
 &= A(\vec{d}; \vec{\mu}, \Sigma)
 \end{aligned}$$

Henceforth, $A(\vec{d}; \vec{\mu}, \Sigma)$ will denote the probability that all of an individual's biomarker levels are below the LODs for a given $\vec{\mu}$ and Σ . For the likelihoods of all the cases of $\mathbf{Q}(\vec{Z})$ in between, $L_k(\vec{\mu}, \Sigma; \vec{Z})$, $k = 2, 3, \dots, 2^p - 1$ (ie, those $\vec{Z}$ composed of both quantified and censored levels), we will extend ideas from explicit results for $P = 2$ to the general p case here (8,15). In that work, the bivariate normal case with $\mathbf{Q}(\vec{Z}) = (1, 0)$ or $(0, 1)$ was shown to lead to a likelihood that is the product of a marginal distribution for the biomarker that is observed and a conditional distribution for the biomarker measured below the LOD. We will rely on the same rationale here but allow the observed and unobserved to be vectors of biomarkers with marginal and conditional distributions that can themselves be multivariate.

For mixed vectors, reorder and partition the observed vector such that $\vec{Z} = (\vec{Z}^{(1)}, \vec{Z}^{(2)})$, where $\vec{Z}^{(1)} = \{W_l; Q(\vec{Z})_l = 1, l = 1, \dots, p\}$ and $\vec{Z}^{(2)} = \{Q(\vec{Z})_l = 0, l = 1, \dots, p\}$. This reordering is achieved by multiplying by a transformation matrix, $\mathbf{C}$, based on $\mathbf{Q}(\vec{Z})$. For example, if $p = 5$ and $\vec{Z} = (., w_2, w_3, ., w_5)$, then based on $\mathbf{Q}(\vec{Z}) = (0, 1, 1, 0, 1)$, we would form the following partition

$$\mathbf{C} \vec{Z} = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix} \begin{pmatrix} . \\ w_2 \\ w_3 \\ . \\ w_5 \end{pmatrix} = \begin{pmatrix} w_2 \\ w_3 \\ w_5 \\ . \\ . \end{pmatrix} = \begin{pmatrix} \vec{Z}^{(1)} \\ \vec{Z}^{(2)} \end{pmatrix}$$

A consequence of this reordering is that $\mathbf{C}\vec{Z} \sim N_p(\mathbf{C}\vec{\mu}, \mathbf{C}\Sigma\mathbf{C}^T)$ and that $\mathbf{C}\vec{\mu} = (\vec{\mu}^{(1)}, \vec{\mu}^{(2)})^T$

and $\mathbf{C}\Sigma\mathbf{C}^T = \begin{bmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{bmatrix}$ can be similarly partitioned, where Σ_{ab} can be a matrix, vector, or scalar depending on the partition. The resultant multivariate normal density is then factored into a marginal density, $g_{(1)}(\vec{W}^{(1)}; \vec{\mu}^{(1)}, \Sigma_{11})$, and a conditional density, $g_{(2|1)}(\vec{W}^{(2)}; \vec{\mu}^{(2)}, \Sigma_Q)$, with $\vec{\mu}^{(2)} = \vec{\mu}^{(2)} + \sum_{21} \sum_{11}^{-1} (\vec{W}^{(1)} - \vec{\mu}^{(1)})$ and $\Sigma_Q = \Sigma_{22} - \sum_{21} \sum_{11}^{-1} \sum_{12}$. Since the observations in $\vec{Z}^{(2)}$ are unobserved but below the LODs, the likelihood for the k th unique $\mathbf{Q}(\vec{Z})$, $k = 2, 3, \dots, 2^p - 1$ is

$$\begin{aligned}
L_k \left(\vec{\mu}, \Sigma; \vec{Z}^{(1)}, \vec{Z}^{(2)} \right) &\propto g_{(1)} \left(\vec{w}^{(1)}; \vec{\mu}^{(1)}, \Sigma_{11} \right) \\
&\times \int_{-\infty}^{\vec{d}^{(2)}} g_{(2|1)} \left(\vec{w}^{(2)}; \vec{\mu}^{(Q)}, \Sigma_Q \right) d\vec{w}^{(2)} \\
&= g_{(1)} \left(\vec{w}^{(1)}; \vec{\mu}^{(1)}, \Sigma_{11} \right) \times \Lambda_{(2|1)} \left(\vec{d}^{(2)}; \vec{\mu}^{(Q)}, \Sigma_Q \right)
\end{aligned}$$

Therefore, for a random sample $\mathbf{Z}$ of independent identically distributed p -dimensional vectors $\vec{Z}_j, j = 1, \dots, n$, following a multivariate normal distribution, the likelihood has the form Maximizing Equation A1 with respect to each element of μ and Σ results in MLEs $\hat{\vec{\mu}}$ and $\hat{\Sigma}$. This must be performed numerically due to the complexity of the likelihood for which a closed form solution does not exist for even the uni-variate case. Also of note is that for the explicit $P = 2$ case, Equation A1 reduces to the likelihood of Lyles et al (8). However, out of necessity, we maintain matrix notation rather than the explicit expansion displayed there.

$$\begin{aligned}
L(\vec{\mu}, \Sigma; \mathbf{Z}) &= \prod_{j=1}^{n_Y} L(\vec{\mu}, \Sigma; \vec{Z}_j) \\
&\propto \prod_{\vec{Q}(\vec{Z}_j)=(1,1,\dots,1)^T} g(\vec{w}_j; \vec{\mu}, \Sigma) \times \dots \\
&\times \prod_{\vec{Q}(\vec{Z}_j)=(1,0,\dots,1,1,0)^T} g_{(1)} \left(\vec{w}_j^{(1)}; \vec{\mu}^{(1)}, \Sigma_{11} \right) \times \Lambda_{(2|1)} \left(\vec{d}^{(2)}; \vec{\mu}^{(Q)}, \Sigma_Q \right) \times \dots \quad (A1) \\
&\times \prod_{\vec{Q}(\vec{Z}_j)=(0,0,\dots,0,0)^T} \Lambda(\vec{d}; \vec{\mu}, \Sigma)
\end{aligned}$$

Fisher Information Matrix

Now that we have estimates $\hat{\vec{\mu}}$ and $\hat{\Sigma}$ based on $\mathbf{Z}$, it is of practical importance to attach some level of certainty to these estimators. The covariance matrices of MLE's of unknown vectors of parameters are evaluated by the inverse of the Fisher information matrices, say I ; that is

$$\begin{aligned}
Cov(\vec{\theta}) &= Cov(\hat{\vec{\mu}}, \hat{\Sigma}) = [I]^{-1} \\
&= \begin{bmatrix} I_{1,1} & \cdots & I_{1,p(p+1)} \\ \vdots & \ddots & \vdots \\ I_{p(p+1),1} & \cdots & I_{p(p+1),p(p+1)} \end{bmatrix}, \quad (A2)
\end{aligned}$$

where $(P = 2)$ and each $I_{ab} = \lim_{n \rightarrow \infty} \frac{1}{n} E \left[\frac{\partial^2 \log(L(\vec{\theta}; \mathbf{Z}))}{\partial \theta_a \partial \theta_b} \right]$. Please note that the rows and columns of Equation A2 correspond to $\vec{\theta} = (\vec{\mu}, \vec{\Sigma})^T = (\mu_1, \dots, \mu_p, \sigma_{11}, \sigma_{12}, \dots, \sigma_{pp})^T$, a vector where the elements of Σ are listed out as a vector and concatenated with $\vec{\mu}$. The development of the I_{ab} , based on the partial derivatives of the log likelihood in Equation A1, was done

with respect to the subvectors and matrices in $\vec{\mu} = (\vec{\mu}^{(1)}, \vec{\mu}^{(2)})^T$ and $\vec{\Sigma} = \begin{bmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{bmatrix}$ in order to facilitate for the general p dimensionality of our data. The expected values of elements I_{ab} are explicitly evaluated where possible and are approximated otherwise.

Estimated AUC for a BLC with CI

Let p biomarkers for individuals with and without a particular disease follow multivariate normal distributions, $\vec{X} \sim N_p(\vec{\mu}_X, \Sigma_X)$ and $\vec{Y} \sim N_p(\vec{\mu}_Y, \Sigma_Y)$, respectively. Then indices formed

from the linear combinations of those biomarkers, $U = \vec{\beta}^T \vec{X} = \sum_{j=1}^p \beta_j X_j$ and

$V = \vec{\beta}^T \vec{Y} = \sum_{j=1}^p \beta_j Y_j$, are univariate normally distributed, $U \sim N_1(\vec{\beta}^T \vec{\mu}_X, \vec{\beta}^T \Sigma_X \vec{\beta})$ and $V \sim N_1(\vec{\beta}^T \vec{\mu}_Y, \vec{\beta}^T \Sigma_Y \vec{\beta})$, respectively. Consequentially, an ROC curve can be formed from U and V that depends on the choice of $\vec{\beta}$. The choice of $\vec{\beta}$ can be equally weighted or an alternative designated by expert knowledge. In this setting, Su and Liu developed the BLC of multiple biomarkers

$$\vec{\beta}_0^T \propto (\vec{\mu}_X - \vec{\mu}_Y)^T (\Sigma_X + \Sigma_Y)^{(-1)}, \quad (A3)$$

leading to

$$AUC_0 = P(U > V) \\ = \Phi \left(\sqrt{(\vec{\mu}_X - \vec{\mu}_Y)^T (\Sigma_X + \Sigma_Y)^{-1} (\vec{\mu}_X - \vec{\mu}_Y)} \right), \quad (A4)$$

which is “best” with respect to maximizing AUC over all real $\vec{\beta}$ (16).

Using samples measured with LOD, $\mathbf{Z}_X$, and $\mathbf{Z}_Y$, we can maximize the likelihood in

Equation A1 for both groups generating MLEs $\hat{\vec{\mu}}_X$, $\hat{\Sigma}_X$, $\hat{\vec{\mu}}_Y$, and $\hat{\Sigma}_Y$. Placing these MLEs in

Equations A3 and A4 results in MLEs $\hat{\vec{\beta}}_0^T$ and $A\hat{U}C_0$. Without loss of generality, we assume that both diseased and nondiseased biomarker levels are affected by the censoring at the same level, $d_{xI} = d_{yI} = d_I$.

When $A\hat{U}C_0$ is based on two normally distributed biomarkers measured without LOD, Reiser and Faraggi suggest creating a CI for AUC_0 through

$\delta = \sqrt{(\vec{\mu}_X - \vec{\mu}_Y)^T (\Sigma_X + \Sigma_Y)^{-1} (\vec{\mu}_X - \vec{\mu}_Y)}$ due to the bounded nature of AUC_0 and the potentially unpredictable convergence of the MLE $A\hat{U}C_0$ to normality around those bounds (17). Here, we assume only the availability of $\mathbf{Z}_X$ and $\mathbf{Z}_Y$. Then, the MLE $\hat{\delta}$, generated similarly to $A\hat{U}C_0$, is still asymptotically distributed $\sqrt{N}(\hat{\delta} - \delta) \sim Normal(0, \sigma_\delta^2)$. The asymptotic variance σ_δ^2 is obtained by

$$\sigma_{\delta}^2 = \lambda^{-1} \left(\frac{\partial \delta}{\partial \vec{\theta}_x} \right)^T \left(Cov \left(\vec{\hat{\theta}}_x \right) \right) \left(\frac{\partial \delta}{\partial \vec{\theta}_x} \right) + (1-\lambda)^{-1} \left(\frac{\partial \delta}{\partial \vec{\theta}_y} \right)^T \left(Cov \left(\vec{\hat{\theta}}_y \right) \right) \left(\frac{\partial \delta}{\partial \vec{\theta}_y} \right), \quad (A5)$$

with $\vec{\theta} = (\vec{\mu}, \vec{\Sigma})^T$ have the form described in Materials and Methods, and $\left(\frac{\partial \delta}{\partial \vec{\theta}} \right)_{ij} = \frac{\partial \delta}{\partial \theta_{ij}}$

being the ij^{th} element of $\left(\frac{\partial \delta}{\partial \vec{\theta}} \right)$, the partial derivative of δ with respect to the vectorized parameters. A $1 - \alpha$ level CI for $A\hat{U}C_0$ can subsequently be calculated by

$\Phi \left(\hat{\delta} \pm z_{\alpha/2} \frac{\sqrt{\sigma_{\delta}^2}}{\sqrt{N}} \right)$. Given $\hat{\vec{\mu}}_x$, $\hat{\vec{\mu}}_y$, $\hat{\Sigma}_x$, and $\hat{\Sigma}_y$ based on $\mathbf{Z}_x$ and $\mathbf{Z}_y$, we use this technique to calculate an approximately $1 - \alpha$ level CI for $A\hat{U}C_0$.

Figure 1.

Relative bias and root mean square error of maximum likelihood estimates of various levels of AUC_0 (0.6 triangles, 0.8 squares) based on simulated multivariate normally distributed data. Censoring denotes limit of detection set at each quartile of the marginal distributions. Sample sizes 50, 100, and 200 correspond to increasing character size. AUC_0 is the area under the ROC curve for the best linear combination.

Figure 2.

QQ normal plots of the log transformed biomarker levels of polychlorinated biphenyls (PCBs) 180, 153, and 188 for women with endometriosis (cases) and without (controls). Horizontal points correspond to data below the limit of detection that are indistinguishable from each other.

Figure 3. Binormal parametric receiver operating characteristic curves for log polychlorinated biphenyl (PCB) 180 (dashed) and the best linear combination of log PCBs 180, 153, and 188.