Quantifying the Effects of Mask Metadata Disclosure and Multiple Releases on the Confidentiality of Geographically Masked Health Data

Dale L. Zimmerman[1] and Claire Pavlik[2]

July 25, 2006

[1]Dale L. Zimmerman is Professor, Department of Statistics and Actuarial Science and Department of Biostatistics, University of Iowa, Iowa City, IA 52242 (E-mail: dzimmer@stat.uiowa.edu; Phone: 319-335-0818; Fax: 319-335-3017), and Affiliate of the Center for Health Policy and Research, College of Public Health, University of Iowa. Please address all correspondence to this author. This author's research was supported by Cooperative Agreement #S-3111 between the Centers for Disease Control and Prevention (CDC) and the Association of Schools of Public Health (ASPH); its contents are the responsibility of the authors and do not necessarily reflect the official views of the CDC or ASPH.

[2]Claire Pavlik is Lecturer, Department of Geography, and Affiliate of the Center for Health Policy and Research, College of Public Health, University of Iowa, Iowa City, IA 52242.

Abstract

The availability of individual-level health data presents opportunities for monitoring the distribution and spread of emergent, acute, and chronic conditions, as well as challenges with respect to maintaining the anonymity of persons with health conditions. Particularly when such data are mapped as point locations, concerns arise regarding the ease with which individual identities may be determined by linking geographic coordinates to digital street networks, then determining residential addresses and, finally, names of occupants at specific addresses. The utility of such datasets must therefore be balanced against the requirements of protecting the confidentiality of individuals whose identities might be revealed through the availability of precise and accurate locational data. Recent literature has pointed towards geographic masking as a means for striking an appropriate balance between data utility and confidentiality. However, questions remain as to whether certain characteristics of the mask (mask metadata) should be disclosed to data users and whether two or more distinct masked versions of the data can be released without breaching confidentiality. In this article, we address these questions by quantifying the extent to which the disclosure of mask metadata and the release of multiple masked versions may affect confidentiality, with a view towards providing guidance to custodians of health datasets. The masks considered include perturbation, areal aggregation, and their combination. Confidentiality is measured by the areas of confidence regions for individuals' locations, which are derived under the probability models governing the masks, conditioned on the disclosed mask metadata.

Key words: Disclosure risk, Locational error, Perturbation, Privacy, Spatial aggregation, Spatial masking.

Introduction

Access to data on the geographic distribution of health conditions is important to public health officials, academic researchers, and the general public. Public health officials need spatially linked health information in order to direct prevention and control activities to areas of need; researchers require access in order to conduct spatial analyses addressing important scientific and public policy questions, many of which may not have been envisioned when the data were originally collected; and ordinary citizens naturally want access to information that is relevant to their own individual health status (e.g. the locations of unusually high rates of cancer and other diseases). However, particularly since the 1996 enactment of HIPAA regulations in the United States, the legal requirements of maintaining the confidentiality of health information have led to heightened concern about disclosure of individual-level health information and an increased interest in developing strategies that permit access. This is particularly so for health data that include the geographic coordinates of affected individuals since the ability of inverse address matching technology to reveal the street address of a domicile at a point location and the names of its residents makes disclosure of individual identities straightforward. Thus, it is widely recognized, within government agencies and other custodians of health data as well as within the research community, that methods are needed for providing access to data with sufficient detail to understand and evaluate the spatial distribution of health conditions, while at the same time sufficiently preserving the anonymity of individuals.

An increasingly important strategy for permitting access to sensitive data while protecting individual identities is to "mask" the data before releasing it to legitimate users. Masking includes but is not limited to the removal or encryption of obvious identifiers such as names, residential addresses, and social security numbers; it may also involve relatively more sophisticated statistical disclosure limitation procedures such as sampling, adding simulated data, grouping, and swapping (Duncan and Pearson, 1991). For spatial data, effective masking requires the modification of geographic coordinates linked to each individual so that inverse

address matching can be circumvented. Two important spatial or geographic masks are random perturbation ("jittering") and areal aggregation; see for example, Armstrong, Rushton, and Zimmerman (1999), French and Wand (2004), Kwan, Casas, and Schmitz (2004), Leitner and Curtis (2004), and Fefferman, O'Neil, and Naumova (2005). The goal of geographically masking a geocoded health dataset is to reduce the potential for identification of affected individuals to acceptably low levels, while at the same time retaining sufficient geographic detail to permit accurate spatial analyses of the data.

Although geographic masking has the potential to strike a satisfactory compromise between the imperatives of confidentiality and geographic detail, two important aspects of the overall problem have heretofore received scant attention. We consider these in the context of a typical situation in which an organization or agency serves as the custodian of a dataset of health care information. As data custodian, it is responsible for maintaining information on occurrences of a specified health condition with temporal and spatial markers, including residential addresses that have been geocoded to geographic coordinates, and for making the data available to appropriate outside users. Upon receiving a request for a dataset from a legitimate user, the custodian will release a geographically masked version of it. But should the custodian also reveal specific information about the geographic masking procedure, i.e. mask metadata, to the user? The disclosure of mask metadata would aid researchers by allowing them, in principle, to determine how definitive the conclusions of their spatial analysis are, or (via power calculations) whether conducting such an analysis is even worthwhile. However, it may also assist a "hacker" whose goal is to identify individuals in the dataset. For example, disclosing that the locations of disease cases have been randomly perturbed, with perturbations sampled randomly from a uniform distribution within a circle of specified radius, may assist a legitimate data user to determine whether any apparent clusters of cases are statistically significant (at, say, the 0.05 level). But it may also help a hacker to breach confidentiality by eliminating from consideration as possible cases a substantial proportion of the background population (specifically, those individuals who do not reside in any of the

circles centered on masked locations). Thus the custodian must carefully consider the type and specificity of information disclosed in mask metadata and its effects on confidentiality.

A second question the data custodian may face is whether to release multiple distinct masked versions of the data. This question arises naturally because the type and extent of geographic masking applied to the data affects the kinds of spatial analyses that may be performed and because different kinds of analyses, corresponding to a range of research goals, may be desired by legitimate users over time. However, a potential consequence of multiple releases is that when combined, they may appreciably enhance a hacker's ability to identify individuals. Therefore, it is essential that the data custodian consider the effects that the release of multiple masked versions may have on confidentiality.

The purpose of this article is to quantify the extent to which confidentiality can be affected by the disclosure of mask metadata and by multiple masked releases, in the hope that such a quantification may provide useful guidance to data custodians. To this end, a quantitative measure of confidentiality or its complement, disclosure risk, must be devised. Many such measures have been devised for general data confidentiality settings; see, for example, Duncan and Lambert (1989). Some measures adapted to geographic masking in particular are proposed by Armstrong, Rushton, and Zimmerman (1999), Ohno-Machado, Silveira, and Vinterbo (2004), and VanWey, Rindfuss, Gutmann, Entwisle, and Balk (2005). Here, we focus on a quantity that purports to measure the risk of what Duncan and Lambert (1989) call identity disclosure, i.e. the ease with which a record in the masked data can be linked to a specific person (henceforth called a "case"). This quantity is the area of a confidence region (of specified coverage probability) for the true location of a case, which is derived under the probability model governing the mask(s), conditioned on the disclosed mask metadata. Because a larger confidence region corresponds to lesser knowledge about the true case location, the larger this measure's value the greater the level of confidentiality (and the smaller the disclosure risk). To a hacker, the practical implications of a larger confidence region are that greater time, effort, and expense would be required to identify an

individual.

To simplify interpretation, our framework for masking and measuring confidentiality is initially area-based rather than population-based; that is, it does not account for variation in the population density within subregions of the study region (or equivalently, it assumes that the density is constant over the study region). However, adapting both the masking protocols and the quantification of confidentiality to a situation with a spatially-varying population density is straightforward, as we will demonstrate.

The remainder of the paper is organized as follows. The next section quantifies confidentiality for some situations in which perturbation masks are employed in all releases, but with different levels of specificity in the disclosed mask metadata. The two sections after that describe modifications that allow for a spatially-varying population density and for different perturbation dispersion parameters across releases. The penultimate section quantifies confidentiality for situations in which one or more of the masks are areal aggregation masks. Major conclusions are discussed in the final section.

Multiple Perturbation Masks

Assume that the data custodian possesses a list of the addresses of k cases or generic health events and plans to release n perturbed versions of these case locations, where $n \geq 1$. Denote the true (x, y) coordinates of these locations by $\boldsymbol{\mu}_1, \boldsymbol{\mu}_2, \ldots, \boldsymbol{\mu}_k$. Suppose that each case is masked by adding to the true location's coordinates a perturbation obtained by randomly sampling from a common probability distribution and that this sampling is performed independently across cases. While many choices exist for this probability distribution, we consider the two that are likely to be used most often in practice: a bivariate normal distribution with equal means 0, equal variances σ^2, and correlation zero; and a uniform distribution on a circular region of radius r centered at 0. We write these two distributions as $\mathrm{N}(\mathbf{0}, \sigma^2\mathbf{I})$ and $\mathrm{U}[C(\mathbf{0}, r)]$, respectively; here and throughout, we use $C(\mathbf{a}, r)$ to represent a circle of radius r centered at a point $\mathbf{a}$. We refer to the former distribution as a circular normal

distribution, and to the latter as a circular uniform distribution. Denote the perturbed locations by $\mathbf{Z}_{ij} = (x_{ij}, y_{ij})'$ for $i = 1, 2, \ldots, k$ and $j = 1, 2, \ldots, n$, and let $\mathbf{Z}$ denote this entire set of locations.

Although the data custodian knows which case each perturbed location corresponds to, this is not necessarily so for the data user. For this knowledge to be conveyed to the data user, generic or encrypted labels identifying the case to which each perturbed location corresponds must be assigned and disclosed with each masked release. Figure 1 illustrates this notion for a simple example in which there are four cases and three masked releases. In the left panel, the case labels for the perturbed locations are disclosed. In the right panel, however, case labels are not disclosed, and attempts to group perturbed locations by case are prone to error; note that simply grouping locations that are "closest together" in any reasonable sense would mistakenly assign one group 2 location to group 3 and one group 3 location to group 2. Because disclosure of case labels affects confidentiality differently than non-disclosure, we consider these two possibilities separately in what follows. Moreover, we assume that the form of the perturbation distribution (circular normal or circular uniform) is disclosed to data users, but that the dispersion parameter (σ^2 or r) may or may not be disclosed.

Case Labels Disclosed

Normal distribution with σ^2 disclosed

As our first scenario we suppose that perturbations are randomly sampled from a $\mathrm{N}(\mathbf{0}, \sigma^2\mathbf{I})$ distribution, where case labels and the value of σ^2 are disclosed to data users. Conveniently, this is a standard grouped-data multivariate normal sampling situation, for which maximum likelihood estimators (MLEs) and confidence regions for the $\boldsymbol{\mu}_i$'s are well-known. From the point of view of a hacker, the likelihood function is

$$L(\boldsymbol{\mu}_1, \ldots, \boldsymbol{\mu}_k | \mathbf{Z}, \sigma^2) \propto \exp[-(1/2\sigma^2) \sum_{i=1}^{k} \sum_{j=1}^{n} (\mathbf{Z}_{ij} - \boldsymbol{\mu}_i)'(\mathbf{Z}_{ij} - \boldsymbol{\mu}_i)]. \tag{1}$$

It follows easily that the MLE of $\boldsymbol{\mu}_i$ is $\hat{\boldsymbol{\mu}}_i = \overline{\mathbf{Z}}_i = \frac{1}{n} \sum_{j=1}^{n} \mathbf{Z}_{ij}$ for $i = 1, \ldots, k$. Moreover, $\overline{\mathbf{Z}}_i \sim$

$\mathrm{N}(\boldsymbol{\mu}_i, \frac{\sigma^2}{n}\mathbf{I})$, by which we obtain the following $100(1-\alpha)\%$ confidence region for $\boldsymbol{\mu}_i$:

$$\{\boldsymbol{\mu}_i : (\overline{\mathbf{Z}}_i - \boldsymbol{\mu}_i)'(\overline{\mathbf{Z}}_i - \boldsymbol{\mu}_i) \leq (\sigma^2/n)\chi^2_{2,\alpha}\} \tag{2}$$

where $\chi^2_{2,\alpha}$ is the $100(1-\alpha)$th percentile of a chi-square distribution with two degrees of freedom. Note that this confidence region is a circle of radius $\sqrt{(\sigma^2/n)\chi^2_{2,\alpha}}$ centered at $\overline{\mathbf{Z}}_i$; thus, the area of the confidence region is

$$\pi\sigma^2\chi^2_{2,\alpha}/n. \tag{3}$$

Expression (3) is useful in two ways. First, it indicates that the level of confidentiality, as measured by the area of a confidence region of fixed confidence level, decreases at the rate of $1/n$ as the number n of masked releases increases. Second, it can provide guidance to the data custodian; if, for instance, a policy existed on the minimal allowable area of a 95% confidence region for each $\boldsymbol{\mu}_i$, then expression (3) could be used to solve for the largest value of n (for fixed σ^2) or the smallest value of σ^2 (for fixed n) that would keep the data custodian in compliance with the policy.

Normal distribution with σ^2 undisclosed

Now consider a setting similar to that of the previous subsection, but in which the value of σ^2 is not disclosed to data users. The hacker's likelihood function here is similar to that given by (1), but reflects the fact that σ^2 is undisclosed and must therefore be regarded as an unknown parameter to be estimated:

$$L(\boldsymbol{\mu}_1, \ldots, \boldsymbol{\mu}_k, \sigma^2|\mathbf{Z}) \propto (\sigma^2)^{-nk} \exp[-(1/2\sigma^2)\sum_{i=1}^{k}\sum_{j=1}^{n}(\mathbf{Z}_{ij} - \boldsymbol{\mu}_i)'(\mathbf{Z}_{ij} - \boldsymbol{\mu}_i)].$$

Nevertheless, this is again equivalent to a standard (but slightly different) grouped-data multivariate normal sampling situation, for which the MLEs of $\boldsymbol{\mu}_i$ $(i = 1, \ldots, k)$ and σ^2 and the theory that leads to a confidence region for each $\boldsymbol{\mu}_i$ are well-known. Note that there must be at least two releases $(n \geq 2)$ for σ^2 to be estimable and for confidence regions to

exist. The MLEs of $\boldsymbol{\mu}_i$ and σ^2 are $\hat{\boldsymbol{\mu}}_i = \overline{\mathbf{Z}}_i$ and

$$\hat{\sigma}^2 = \frac{1}{2nk} \sum_{i=1}^{k} \sum_{j=1}^{n} (\mathbf{Z}_{ij} - \overline{\mathbf{Z}}_i)'(\mathbf{Z}_{ij} - \overline{\mathbf{Z}}_i),$$

respectively. Now we have that

$$n(\overline{\mathbf{Z}}_i - \boldsymbol{\mu}_i)'(\overline{\mathbf{Z}}_i - \boldsymbol{\mu}_i)/\sigma^2 \sim \chi^2_2 \quad \text{and} \quad 2nk\hat{\sigma}^2/\sigma^2 \sim \chi^2_{2k(n-1)},$$

and that these two quantities are distributed independently. Thus, a $100(1-\alpha)\%$ confidence region for $\boldsymbol{\mu}_i$ is given by

$$\left\{ \boldsymbol{\mu}_i : (\overline{\mathbf{Z}}_i - \boldsymbol{\mu}_i)'(\overline{\mathbf{Z}}_i - \boldsymbol{\mu}_i) \leq \frac{2\hat{\sigma}^2}{n-1} F_{2,2k(n-1),\alpha} \right\} \tag{4}$$

where $F_{2,2k(n-1),\alpha}$ is the $100(1-\alpha)$th percentile of an F distribution with 2 and $2k(n-1)$ degrees of freedom. This region is a circle centered at $\overline{\mathbf{Z}}_i$, with expected area

$$\frac{2\pi}{n-1} F_{2,2k(n-1),\alpha} E(\hat{\sigma}^2) = \frac{2\pi\sigma^2}{n} F_{2,2k(n-1),\alpha}. \tag{5}$$

Note that (5) is an expected area, as the area of (4) is a random variable. Also observe that (5), like (3), is constant across cases and decreases at the rate of $1/n$; furthermore, it is somewhat larger than (3).

The ratio of (3) to (5), i.e. $\chi^2_{2,\alpha}/2F_{2,2k(n-1),\alpha}$, can be interpreted as a measure of the degradation in confidentiality due to disclosure of the perturbation variance σ^2: the smaller the ratio, the greater the degradation. Table 1 gives this ratio for selected values of k, n, and α. These results show that the disclosure of σ^2 affects confidentiality less as k and n increase. This is not surprising in light of the fact that the masked data themselves provide information about σ^2, which can be used to construct an ever more precise estimate as k and n increase.

Circular uniform distribution with radius disclosed

Now suppose that perturbed locations are obtained by randomly sampling from a $\mathrm{U}[C(\mathbf{0}, r)]$ distribution, where case labels and the value of r are disclosed to data users. In this situation the hacker's likelihood function can be written as

$$L(\boldsymbol{\mu}_1, \ldots, \boldsymbol{\mu}_k | \mathbf{Z}, r) = \prod_{i=1}^{k} \ell(\boldsymbol{\mu}_i | \mathbf{Z}_i, r) \tag{6}$$

where $\mathbf{Z}_i$ denotes the set of perturbed locations corresponding to case i and

$$\ell(\boldsymbol{\mu}_i | \mathbf{Z}_i, r) = \begin{cases} (1/\pi r^2)^n, & \text{if } (\mathbf{Z}_{ij} - \boldsymbol{\mu}_i)'(\mathbf{Z}_{ij} - \boldsymbol{\mu}_i) < r \text{ for all } j = 1, \ldots, n \\ 0, & \text{otherwise.} \end{cases}$$

Because the unknown parameters appear in only one term each in the product in (6), inference for each true location can be treated separately from the others. In particular, to obtain the MLE of $\boldsymbol{\mu}_i$ it suffices to maximize each $\ell(\boldsymbol{\mu}_i | \mathbf{Z}_i, r)$ separately. The MLE of $\boldsymbol{\mu}_i$ so obtained is not uniquely determined, but is given by any point in the intersection of the circles $C(\mathbf{Z}_{ij}, r)$ across j, i.e.

$$\hat{\boldsymbol{\mu}}_i \in \cap_{j=1}^{n} C(\mathbf{Z}_{ij}, r). \tag{7}$$

Now, for each j we have

$$(\mathbf{Z}_{ij} - \boldsymbol{\mu}_i)'(\mathbf{Z}_{ij} - \boldsymbol{\mu}_i)/r^2 \sim \mathrm{U}(0, 1),$$

and thus

$$\max_j [(\mathbf{Z}_{ij} - \boldsymbol{\mu}_i)'(\mathbf{Z}_{ij} - \boldsymbol{\mu}_i)/r^2] \sim \mathrm{Beta}(n, 1).$$

It follows that a $100(1-\alpha)\%$ confidence region for $\boldsymbol{\mu}_i$ is given by

$$\{\boldsymbol{\mu}_i : \max_j [(\mathbf{Z}_{ij} - \boldsymbol{\mu}_i)'(\mathbf{Z}_{ij} - \boldsymbol{\mu}_i)/r^2] \leq (1-\alpha)^{1/n}\},$$

or equivalently by

$$\{\boldsymbol{\mu}_i : \boldsymbol{\mu}_i \in \cap_{j=1}^{n} C(\mathbf{Z}_{ij}, r(1-\alpha)^{1/2n}\}. \tag{8}$$

This is again an intersection (across j) of circles centered at the $\mathbf{Z}_{ij}$'s, but the common radius of these circles is smaller than that which defines a MLE. It is worth noting that if $n \geq 2$: (a) the shape and area of this confidence region vary with i; and (b) the confidence region is likely to be the empty set, and hence of no use to a hacker, when $1 - \alpha$ is sufficiently small.

When $n = 1$, (8) is simply a circle of area

$$\pi r^2(1 - \alpha). \tag{9}$$

When $n = 2$, (8) is the intersection of two circles with common radius, hence simple geometric considerations yield its area as

$$\begin{cases} r^2(1-\alpha)^{1/2}(q_i - \sin q_i), & \text{if } d_i \leq 2r(1-\alpha)^{1/4} \\ 0, & \text{otherwise,} \end{cases} \tag{10}$$

where $q_i = 2\cos^{-1}\{d_i/[2r(1-\alpha)^{1/4}]\}$ and $d_i = [(\mathbf{Z}_{i1} - \mathbf{Z}_{i2})'(\mathbf{Z}_{i1} - \mathbf{Z}_{i2})]^{1/2}$. This area is a random variable with an analytically intractable expectation, so we estimate the expected area by Monte Carlo simulation. Table 2 gives (9), the exact area when $n = 1$, and the ratio of estimated expected area (10) when $n = 2$ to (9), for selected values of r and α. When $n \geq 3$, an explicit formula for the area of (8), or for its expectation, is intractable. Nevertheless, for a fixed confidence level the effect that increasing n has on this area can be evaluated computationally. Additional rows of Table 2 display the results of one such evaluation. A single case was assumed to occur at (0,0), and n independent U$[C(\mathbf{0}, r)]$ perturbations of this case were obtained for selected values of n and r. A very fine grid (1000×1000) was laid over the study region and the number of grid points contained in (8) (with α taken to be 0.01 or 0.05) was counted. This process was repeated 1000 times for each combination of n, r and α, and the average of these 1000 counts was taken as an estimate of the expected area of (8) for that combination. The rows of Table 2 corresponding to $n \geq 3$, like those for $n = 2$, give the ratio of estimated expected area of (8) to $\pi r^2(1 - \alpha)$. Table 2 as a whole reveals that doubling r from 0.05 to 0.10 approximately quadruples the expected area and doubling

n results in a 54-68% reduction in expected area (over the given range of n). Recalling that in the circular normal case, doubling σ^2 exactly quadruples the area of the confidence region and doubling n exactly halves the area of the confidence region, we see that the effect of multiple releases is similar, but perhaps slightly more detrimental to confidentiality, when the perturbation distribution is circular uniform than when it is normal.

Circular uniform distribution with radius undisclosed

If r is not disclosed, the hacker's likelihood function is slightly different from (6) and can be written as

$$L(\boldsymbol{\mu}_1, \ldots, \boldsymbol{\mu}_k, r|\mathbf{Z}) = \prod_{i=1}^{k} \ell(\boldsymbol{\mu}_i, r|\mathbf{Z}_i) \tag{11}$$

where

$$\ell(\boldsymbol{\mu}_i, r|\mathbf{Z}_i) = \begin{cases} (1/\pi r^2)^n, & \text{if } (\mathbf{Z}_{ij} - \boldsymbol{\mu}_i)'(\mathbf{Z}_{ij} - \boldsymbol{\mu}_i) < r \text{ for all } j = 1, \ldots, n \\ 0, & \text{otherwise.} \end{cases}$$

Because the common radius r must be estimated in this situation, each case cannot be treated separately. From (11) we see that obtaining the MLEs of $\boldsymbol{\mu}_1, \ldots, \boldsymbol{\mu}_k$ and r amounts to determining values $\hat{\boldsymbol{\mu}}_1, \ldots, \hat{\boldsymbol{\mu}}_k$ and $\hat{r}$ such that the k circles centered at $\hat{\boldsymbol{\mu}}_1, \ldots, \hat{\boldsymbol{\mu}}_k$ with common radius $\hat{r}$ will each cover its corresponding group of perturbed locations, with $\hat{r}$ as small as possible. When $k = 1$ this is a classical problem in computational geometry known as the "smallest enclosing circle problem." There is no closed-form solution to this problem, but several ingenious numerical algorithms have been devised; see, for example, Preparata and Shamos (1988). In our situation $k > 1$, hence a few additional steps are required to obtain the MLEs. First, the smallest enclosing circle must be obtained for each group (case) separately. Let m denote the index corresponding to the group whose smallest enclosing circle has the largest radius among groups, and let r_m denote this radius. Then, $\hat{r} = r_m$, and the MLE of $\boldsymbol{\mu}_m$ is the unique point at the center of the circle that encloses group m. The MLE of each remaining $\boldsymbol{\mu}_i$ is given by any point in the intersection of the circles $C(\mathbf{Z}_{ij}, \hat{r})$ across j, i.e., $\hat{\boldsymbol{\mu}}_i \in \cap_{j=1}^{n} C(\mathbf{Z}_{ij}, \hat{r})$.

Recall that in the context of a circular uniform distribution with known radius, we were able to derive a closed-form expression, namely (8), for a $100(1-\alpha)\%$ confidence region for each $\boldsymbol{\mu}_i$. For the present situation in which r is unknown, however, derivation of a confidence region is intractable. Nevertheless, the form of (8) suggests that a hacker might consider using

$$\{\boldsymbol{\mu}_i : \boldsymbol{\mu}_i \in \cap_{j=1}^{n} C(\mathbf{Z}_{ij}, \hat{r}(1-\alpha)^{1/2n}\} \tag{12}$$

as an *ad hoc* "search region" (of unknown coverage probability) for $\boldsymbol{\mu}_i$. Observe that this region is identical to (8) except that $\hat{r}$ appears in place of r. Monte Carlo simulation can be used to study how the area and coverage probability of (12) compare to those of (8), and thus how the non-disclosure of r affects confidentiality. In particular, the area of (12) can be estimated with little difficulty when $n = 2$; for a given combination of k and r we simply need to obtain (12) for many simulated datasets drawn from the distribution specified by (11), use (10) to compute the area of each, and then average these areas over simulations. The coverage probability of (12) can be estimated by the proportion of simulation×case combinations for which (12) contains the true μ_i.

Results for the area of the largest (across cases) search region so obtained and its estimated coverage probability are given in Table 3. Observe, for example, that the area of the 99% confidence region for an arbitrary case when r is known and the expected area of the largest nominal 99% search region when $k = 10$ are about equal, but that the coverage probability of the latter is only about 57%. This, together with the other results of Table 3, clearly demonstrate that the non-disclosure of r has a positive effect on confidentiality. Therefore, the data custodian should be loathe to disclose it, at least when n and k are small.

Case Labels Undisclosed

In the situations considered so far, it has been assumed that the data custodian attaches generic case labels to the perturbed locations, so that data users can correctly group together, across releases, all locations corresponding to the same case. When this is not so,

computation of the MLE is much less straightforward, for generally one must then maximize the likelihood over all possible groupings. The number of ways to assign the perturbed locations to k groups of size n is $(k!)^{n-1}$, which grows very quickly as n increases. For simplicity, therefore, we consider the case $n = 2$ only. In this case the locations from each of the two masked releases can be assigned labels from $\{1, \ldots, k\}$, but there is no guarantee that the second labeling correctly pairs locations corresponding to the same true case. Let $\boldsymbol{\pi} = (\pi_1, \ldots, \pi_k)'$ denote the $k \times 1$ vector of true case labels for the second release (in terms of the labeling used for the first release), which is merely a permutation of the integers $1, \ldots, k$. Now, from the perspective of a hacker, $\boldsymbol{\pi}$ is an unknown parameter taking on one of $k!$ possible values, which must be estimated together with $\boldsymbol{\mu}_1, \ldots, \boldsymbol{\mu}_k$ and any other unknown model parameters. This lack of knowledge of $\boldsymbol{\pi}$ affects the hacker's inference for $\boldsymbol{\mu}_1, \ldots, \boldsymbol{\mu}_k$.

For the circular normal setting with σ^2 disclosed, the hacker's likelihood function is

$$L(\boldsymbol{\pi}, \boldsymbol{\mu}_1, \ldots, \boldsymbol{\mu}_k | \mathbf{Z}, \sigma^2) \propto \exp\{-(1/2\sigma^2) \sum_{i=1}^{k} [(\mathbf{Z}_{i1} - \boldsymbol{\mu}_i)'(\mathbf{Z}_{i1} - \boldsymbol{\mu}_i) + (\mathbf{Z}_{\pi_i 2} - \boldsymbol{\mu}_i)'(\mathbf{Z}_{\pi_i 2} - \boldsymbol{\mu}_i)]\}. \tag{13}$$

Let $\overline{\mathbf{Z}}_{i\pi} = \frac{1}{2}(\mathbf{Z}_{i1} + \mathbf{Z}_{\pi_i 2})$. Careful analysis of (13) reveals that the MLE of $\boldsymbol{\pi}$ is given by the assignment scheme that minimizes the within-group variation, i.e.,

$$\hat{\boldsymbol{\pi}} = \text{argmin} \sum_{i=1}^{k} [(\mathbf{Z}_{i1} - \overline{\mathbf{Z}}_{i\pi})'(\mathbf{Z}_{i1} - \overline{\mathbf{Z}}_{i\pi}) + (\mathbf{Z}_{\pi_i 2} - \overline{\mathbf{Z}}_{i\pi})'(\mathbf{Z}_{\pi_i 2} - \overline{\mathbf{Z}}_{i\pi})],$$

and the MLE of $\boldsymbol{\mu}_i$ is then $\hat{\boldsymbol{\mu}}_i = \overline{\mathbf{Z}}_{i\hat{\boldsymbol{\pi}}}$. Derivation of a $100(1-\alpha)\%$ confidence region for $\boldsymbol{\mu}_i$ appears to be intractable, but by analogy with (2) it would be natural for a data user to use

$$\{\boldsymbol{\mu}_i : (\overline{\mathbf{Z}}_{i\hat{\boldsymbol{\pi}}} - \boldsymbol{\mu}_i)'(\overline{\mathbf{Z}}_{i\hat{\boldsymbol{\pi}}} - \boldsymbol{\mu}_i) \leq (\sigma^2/2)\chi^2_{2,\alpha}\} \tag{14}$$

as an *ad hoc* search region (of unknown coverage probability) for $\boldsymbol{\mu}_i$.

Table 4 gives empirical coverage probabilities (obtained via Monte Carlo simulation) for (14) corresponding to the situation depicted in Figure 2, in which there are $k = 5$ cases distributed within a square of side 2.0. Three values of σ were considered. For $\sigma = 0.05$,

not one of the 10,000 simulations yielded a misclassification of cases to groups, and the empirical coverage probabilities were completely consistent with nominal levels. For $\sigma = 0.20$, however, approximately 23% of the simulations resulted in misclassifications, all of which involved cases 4 and 5 only. The coverage probabilities corresponding to cases 4 and 5 are slightly reduced, to less than 0.97 for the nominal 99% region and to less than 0.90 for the nominal 95% region. For $\sigma = 1.00$, approximately 83% of the simulations resulted in misclassifications, and this time only 16% of them involved interchanges of cases 4 and 5. The actual coverage probabilities of the nominal 99% and 95% search regions were roughly 0.93 and 0.84, regardless of case. Qualitatively similar results were obtained for other choices of k and other spatial configurations of case locations. Thus, it appears that the effect of not disclosing the case labels is to modestly reduce the actual coverage probability of region (14), with the magnitude of the reduction depending on the magnitude of σ relative to the distance between cases.

For the remaining situations we considered in the previous section, similar modifications to account for undisclosed case labels can be made to inference procedures for $\boldsymbol{\mu}_1, \ldots, \boldsymbol{\mu}_k$. For the circular normal setting with undisclosed σ^2, the MLEs of $\boldsymbol{\pi}$ and $\boldsymbol{\mu}_1, \ldots, \boldsymbol{\mu}_k$ are identical to those just described, and the MLE of σ^2 is given by

$$\hat{\sigma}^2 = (1/4k) \sum_{i=1}^{k} [(\mathbf{Z}_{i1} - \overline{\mathbf{Z}}_{i\hat{\boldsymbol{\pi}}})'(\mathbf{Z}_{i1} - \overline{\mathbf{Z}}_{i\hat{\boldsymbol{\pi}}}) + (\mathbf{Z}_{\hat{\boldsymbol{\pi}}_i 2} - \overline{\mathbf{Z}}_{i\hat{\boldsymbol{\pi}}})'(\mathbf{Z}_{\hat{\boldsymbol{\pi}}_i 2} - \overline{\mathbf{Z}}_{i\hat{\boldsymbol{\pi}}})].$$

In the circular uniform setting with r disclosed, all points belonging to nonempty sets of form (7) for some grouping scheme will contribute the same term to the likelihood, namely $(1/\pi r^2)^2$, so any permutation of $\{1, \ldots, k\}$ yielding k such sets will be an MLE of $\boldsymbol{\pi}$, and any points within such sets will be MLEs of their respective $\boldsymbol{\mu}_i$. When r is not disclosed, the MLE of $\boldsymbol{\pi}$ is the permutation that results in the smallest value of $\hat{r}$ obtained by the procedure described previously, and $\hat{\boldsymbol{\mu}}_1, \ldots, \hat{\boldsymbol{\mu}}_k$ may also be obtained as described there. Consequently, a hacker might attempt to construct search regions analogous to (14) in such situations.

Accounting for Population Density

Our framework for masking and measuring confidentiality has not, to this point, incorporated any aspect of variation in population density into the perturbation distribution or the confidence region's construction. As a consequence, in the circular normal situation for example, the areas of confidence regions for case locations were constant across cases. It is rather simple, however, to account for population density in the perturbation model and thereby obtain confidence or search regions whose areas may vary across cases. A data custodian might view accounting for population density as desirable for various purposes. In particular, it would offer a means of equalizing disclosure risk across all cases, which would be useful if a hacker, given sufficient information to construct confidence regions for two cases, requires less effort to breach confidentiality in the region that has the smaller population.

The key to equalizing disclosure risk across cases is to use a different dispersion parameter for each case's perturbation distribution, specifically one which is inversely related to the population size in the vicinity of the case's true location. This could be implemented in the circular uniform situation by choosing the case-specific dispersion parameters $\{r_i : i = 1, \ldots, k\}$ to be such that the same number of individuals in the background population reside in each $C(\boldsymbol{\mu}_i, r_i)$. In the circular normal situation, case-specific variances $\{\sigma_i^2 : i = 1, \ldots, k\}$ might be chosen in such a way that a given, say 95%, probability contour centered at the true location contains the same number of individuals from the background population.

Whatever scheme the data custodian uses to choose the dispersion parameters, expressions for MLEs, confidence regions, and areas of confidence regions that might be computed by a data hacker are relatively straightforward generalizations of expressions we have given previously. For simplicity, suppose that case labels are disclosed. Then, in the circular normal case with variances disclosed, the $100(1-\alpha)\%$ confidence region for $\boldsymbol{\mu}_i$ and its area are given by expressions identical to (2) and (3), except that σ_i^2 replaces σ^2. Likewise, in the circular uniform case with r disclosed, the $100(1-\alpha)\%$ confidence region for $\boldsymbol{\mu}_i$ and its area when $n = 1$ and $n = 2$ are given by expressions identical to (8), (9), and (10), except that

r_i is substituted for r. Now, if the dispersion parameters are not disclosed, then (provided $n \geq 2$) the $100(1-\alpha)\%$ confidence region for $\boldsymbol{\mu}_i$ and its area in the circular normal case are given by expressions identical to (4) and (5), except that

$$\hat{\sigma}_i^2 = \frac{1}{2n} \sum_{j=1}^{n} (\mathbf{Z}_{ij} - \overline{\mathbf{Z}}_i)'(\mathbf{Z}_{ij} - \overline{\mathbf{Z}}_i)$$

is substituted for $\hat{\sigma}^2$, and the second degree of freedom parameter in the F percentile changes from $2k(n-1)$ to $2(n-1)$. As for the circular uniform case, no explicit confidence region is available, but one might expect that a hacker would use a search region akin to (12) with $\hat{r}_i$ substituted for $\hat{r}$, where $\hat{r}_i$ is obtained by solving the smallest enclosing circle problem separately for case i.

For a masking protocol that directly accounts for population density, a data custodian should measure the effects of multiple releases and non-disclosure of dispersion parameters on confidentiality by the population size (or its expectation), rather than the area (or its expectation), of each confidence or search region.

Varying the Dispersion Parameter Across Releases

Occasionally, a data custodian may wish to vary the dispersion parameter of the perturbation distribution across releases. This could happen, for instance, if there was public concern about the possibility of cancer clusters in a certain region, but the custodian had initially released the data using a dispersion parameter so large that researchers had little statistical power to detect any clusters. A second release using a smaller dispersion parameter would confer greater power upon most tests for clusters, and could thus yield more definitive conclusions.

Suppose, then, that the perturbation dispersion parameters of the circular normal and circular uniform distributions are permitted to vary across releases. Suppose further that case labels are disclosed. Then, in the circular normal situation in which release-specific

variances $\{\sigma^2_{(j)} : j = 1, \ldots, n\}$ are disclosed, the MLE of $\boldsymbol{\mu}_i$ is

$$\hat{\boldsymbol{\mu}}_i = \left(\sum_{j=1}^{n} \frac{1}{\sigma^2_{(j)}} \right)^{-1} \sum_{j=1}^{n} \frac{1}{\sigma^2_{(j)}} \mathbf{Z}_{ij},$$

and a $100(1-\alpha)\%$ confidence region for $\boldsymbol{\mu}_i$ is given by

$$\{\boldsymbol{\mu}_i : (\hat{\boldsymbol{\mu}}_i - \boldsymbol{\mu}_i)'(\hat{\boldsymbol{\mu}}_i - \boldsymbol{\mu}_i) \leq \left(\sum_{j=1}^{n} \frac{1}{\sigma^2_{(j)}} \right)^{-1} \chi^2_{2,\alpha}\}. \tag{15}$$

For the corresponding circular uniform situation in which release-specific radii $\{r_{(j)} : j = 1, \ldots, n\}$ are disclosed, a MLE and $100(1-\alpha)\%$ confidence region for $\boldsymbol{\mu}_i$ are given by expressions identical to (7) and (8) except that $r_{(j)}$ replaces r. Thus, both of these are intersections of circles of varying, rather than common, sizes.

If the dispersion parameters are release-specific but not disclosed, the derivation of MLEs and confidence regions is considerably more complicated. We will report results elsewhere for these situations, as well as for situations in which the dispersion parameters are simultaneously release-specific and case-specific (e.g., to account for population density).

In general, the effect of reducing the perturbation dispersion parameter for one or more releases is to reduce the area of the confidence region for $\boldsymbol{\mu}_i$ below what it would be if all releases used the largest dispersion. For example, in the particular case of two releases and a circular normal perturbation distribution, if the variance of the second release is $1/\nu$ as large as that of the first release, then the confidence region given by (15) has area $1/(\nu+1)$ as large as the area of the confidence region obtained using only the data from the first release.

Combining Perturbation with Aggregation

Often, data custodians choose to release an areal aggregation of data, such as a map or table displaying cases or case counts by census tract or county. We now consider situations in which the data custodian combines this type of aggregated data release together with releases using perturbations of point locations.

If all of the data releases are aggregation masks, the effects of multiple releases are transparent and depend on whether the aggregations are spatially nested or crossed (spatially misaligned). In the nested case, e.g. aggregations to U.S. census block groups and tracts, the impact of masking on confidentiality is determined entirely by how small the population sizes are for the areal units used in the smallest-scale aggregation mask. In the crossed case, e.g. aggregations to census tracts and zip codes, population sizes in the areas of intersection across masked releases determine the extent to which confidentiality is affected.

If a single aggregation mask is combined with one or more perturbation masks, it is clear that the confidentiality of only those cases near the boundary of each areal unit will be affected significantly. Moreoever, the aggregation mask will have relatively little effect on confidentiality if the scale of perturbation is small relative to the dimensions of the areal units. To illustrate this numerically, consider a study area comprising a rectangular array of square areal units, each of side a, and suppose that the background population, as well as the subpopulation of cases, is distributed independently and uniformly over this study area. Suppose further that two masked versions are released: one which perturbs case locations independently according to a uniform distribution on a square of side $2r$ (with $r \leq a/2$) centered at their respective actual locations, and another which aggregates case locations to their respective areal units. For simplicity, we assume that r and the case labels are disclosed, and we consider only those cases whose actual locations lie in an arbitrary interior areal unit S; observe that after perturbation such cases must lie in another square, say T, of side $a + 2r$ containing S (see Figure 3).

Now let $\boldsymbol{\mu} = (u, v)$ be the true coordinates of an arbitrary case in S, let $\mathbf{z} = (x, y)$ be the corresponding perturbed location, and let $B(\mathbf{z}, r)$ be a square of side $2r$ centered at $\mathbf{z}$. From the hacker's perspective, the region which must contain $\boldsymbol{\mu}$, given only the information available from the two masked datasets (namely, that $\boldsymbol{\mu} \in S$ and that its corresponding perturbed location is $\mathbf{z}$), is the rectangle $B(\mathbf{z}, r) \cap S$. Let $A(\mathbf{z}, r)$ be the area of this rectangle. Note that $0 < A(\mathbf{z}, r) \leq 4r^2$; furthermore, $A(\mathbf{z}, r)$ is close to zero if and only if $\mathbf{z}$ is close

to the boundary of T, and the upper bound $4r^2$ is attained if and only if $\mathbf{z}$ is in S and is no closer than r to the boundary of S. Define $\gamma \equiv E\{A(\mathbf{z}, r)\}/(4r^2)$, where the expectation is taken over the distribution of $\mathbf{z}$. We have that $0 < \gamma < 1$, and the closer that γ is to 0 the greater the effect of the aggregation mask on confidentiality relative to the perturbation mask.

To obtain an expression for γ in terms of a and r, let $g(\boldsymbol{\mu})$ and $f(\mathbf{z})$ denote the probability density functions of $\boldsymbol{\mu}$ and $\mathbf{z}$, respectively, and let $h(\mathbf{z}|\boldsymbol{\mu})$ denote the conditional density of $\mathbf{z}$ given $\boldsymbol{\mu}$. Then $f(\mathbf{z})$ is given by

$$\begin{aligned} f(\mathbf{z}) \,=\, f(x,y) &= \int_{B(\mathbf{z},r)\cap S} g(u,v)h(x,y|u,v)\,du\,dv \\ &= \int_{B(\mathbf{z},r)\cap S} \frac{1}{a^2}\cdot\frac{1}{4r^2}\,du\,dv \\ &= \frac{1}{4a^2r^2}A(\mathbf{z},r), \text{ for } \mathbf{z}\in T. \end{aligned}$$

Thus, $\gamma = a^{-2}\int_T A^2(\mathbf{z}, r)\,d\mathbf{z}$. Now, by partitioning T into suitable subregions and exploiting various symmetries, we find, after tedious but straightforward calculations, that $\gamma = (3a - 2r)^2/(9a^2)$. Thus, as expected, γ tends to 1 as a gets large relative to r. Values of γ for selected values of a/r are as follows:

a/r	2	3	4	6	10	20	50
γ	0.444	0.605	0.694	0.790	0.871	0.934	0.974

These results show that if both masked versions of the data are released and if, for example, $a/r = 6$, then the average effort (across cases) a data hacker would have to expend to identify an individual case is 79% of what it would be if only the perturbed data were released. For a case whose masked location is close to a boundary, the required effort could be significantly less.

Conclusions

In this article, we have examined, through a combination of statistical theory and simulation, how confidentiality is affected by the disclosure of mask metadata and the release of

multiple masked versions. To illustrate our approach, we considered situations in which a data custodian perturbs generically labeled point locations using normal bivariate and circular uniform spatial masks, then either conceals or discloses information about the dispersion parameters of the mask(s). We used our results for situations without mask metadata disclosure to generate ratios for examining the reduction in confidentiality that occurs when mask metadata are revealed. As anticipated, disclosing mask metadata decreased the area (or expected area) of confidence regions of fixed coverage probability, and increased the coverage probability of confidence regions of fixed area, with the strongest effects observed for lower numbers of releases and smaller numbers of cases. These effects are intuitively reasonable: given a specific number of cases, increasing the number of releases means that the releases themselves, bereft of mask metadata, will provide additional information on the mask parameter; also, given a specific number of masked data releases, higher numbers of cases also provides additional information about mask parameters. Thus for both higher numbers of masked data releases and higher case counts, the impact of revealing mask metadata is lower. We also considered the relative impact of disclosure versus non-disclosure of case labels, illustrating the results of a specific situation with five cases and two data releases. Our results in this and other instances indicate that not disclosing case labels modestly reduces the coverage probability of a confidence region, with the impact of the reduction related to the relative sizes of the mask's dispersion parameter(s) and the distances between individual cases.

In addition to multiple perturbation masks, we considered the combination of a release of a spatially aggregated dataset and one or more perturbed releases of point data. This combination is quite likely to occur in practice, given the tendency for data custodians to release datasets aggregated to recognizable regions, such as counties or states, for public information, while also releasing one or more masked point datasets for analytical use by researchers. In such cases, a key consideration for the masked datasets is the relative size of the perturbations in relation to the areal units used for aggregation. Particularly when

the areal units are not much larger than the uniform region(s) used for perturbation, the information revealed may shrink a hacker's search regions significantly, thus dramatically compromising confidentiality, especially for those cases located near the boundaries of the areal units.

Although we obtained numerical results for several specific masking scenarios and described how our methodology could be modified to obtain results for other scenarios (namely those that involve release-specific dispersion parameters or account for variation in population density), it is still true that many practical masking scenarios facing data custodians will not match any of the ones we have illustrated in all details. Therefore, we anticipate that data custodians will need to further modify our methodology to fit their particular scenarios.

References

Armstrong, M.P., G. Rushton, and D.L. Zimmerman (1999). "Geographically Masking Health Data to Preserve Confidentiality." *Statistics in Medicine* 18, 497-525.

Duncan, G. and D. Lambert (1989). "The Risk of Disclosure for Microdata." *Journal of Business and Economic Statistics* 7, 207-17.

Duncan, G. and R.W. Pearson (1991). "Enhancing Access to Microdata while Protecting Confidentiality: Prospects for the Future." *Statistical Science* 6, 219-239.

Fefferman, N.H., E.A. O'Neil, and E.N. Naumova (2005). "Confidentiality and Confidence: Is Data Aggregation a Means to Achieve Both?" *Journal of Public Health Policy* 26, 430-449.

French, J.L. and M.P. Wand (2004). "Generalized Additive Models for Cancer Mapping with Incomplete Covariates." *Biostatistics* 5, 177-91.

Kwan, M.P., I. Casas, and B.C. Schmitz (2004). "Protection of Geoprivacy and Accuracy of Spatial Information: How Effective are Geographical Masks?" *Cartographica* 39, 15-28.

Leitner, M. and A. Curtis (2004). "Cartographic Guidelines for Geographically Masking the Locations of Confidential Point Data." *Cartographic Perspectives* 49, 8-25.

Ohno-Machado, L., P.S.P. Silveira, and S. Vinterbo (2004). "Protecting Patient Privacy by Quantifiable Control of Disclosures in Disseminated Databases." *International Journal of Medical Informatics* 73, 599-606.

Preparata, F.P. and M.I. Shamos (1988). *Computational Geometry: An Introduction*. New York: Springer-Verlag.

VanWey, L.K., R.R. Rindfuss, M.P. Gutmann, B. Entwisle, and D.L. Balk (2005). "Confidentiality and Spatial Explicit Data: Concerns and Challenges." *Proceedings of the National Academy of Sciences of the United States of America* 102, 15337-42.

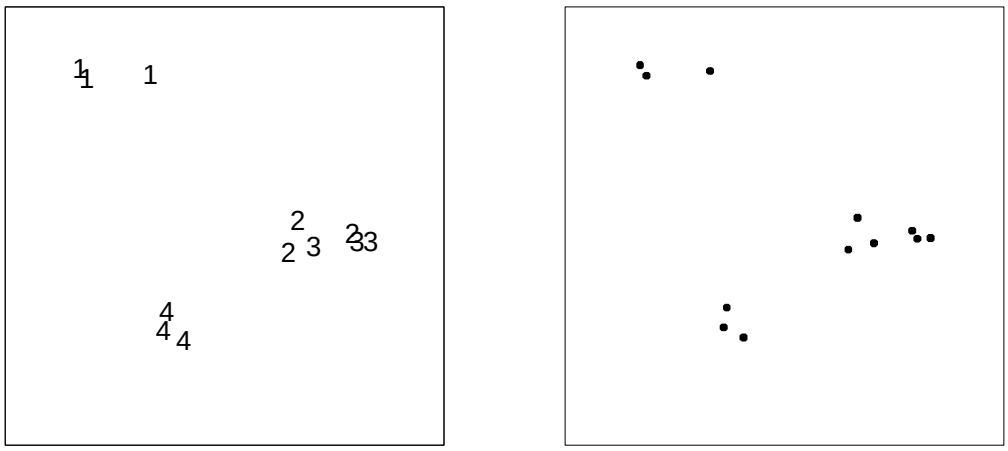

Figure 1

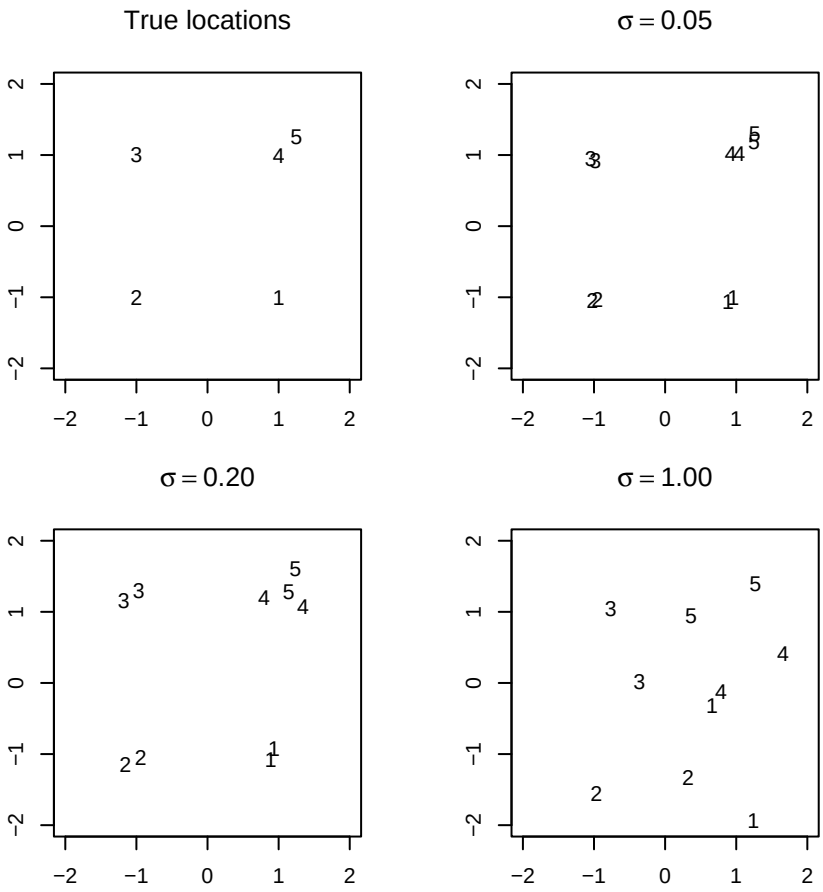

Figure 2

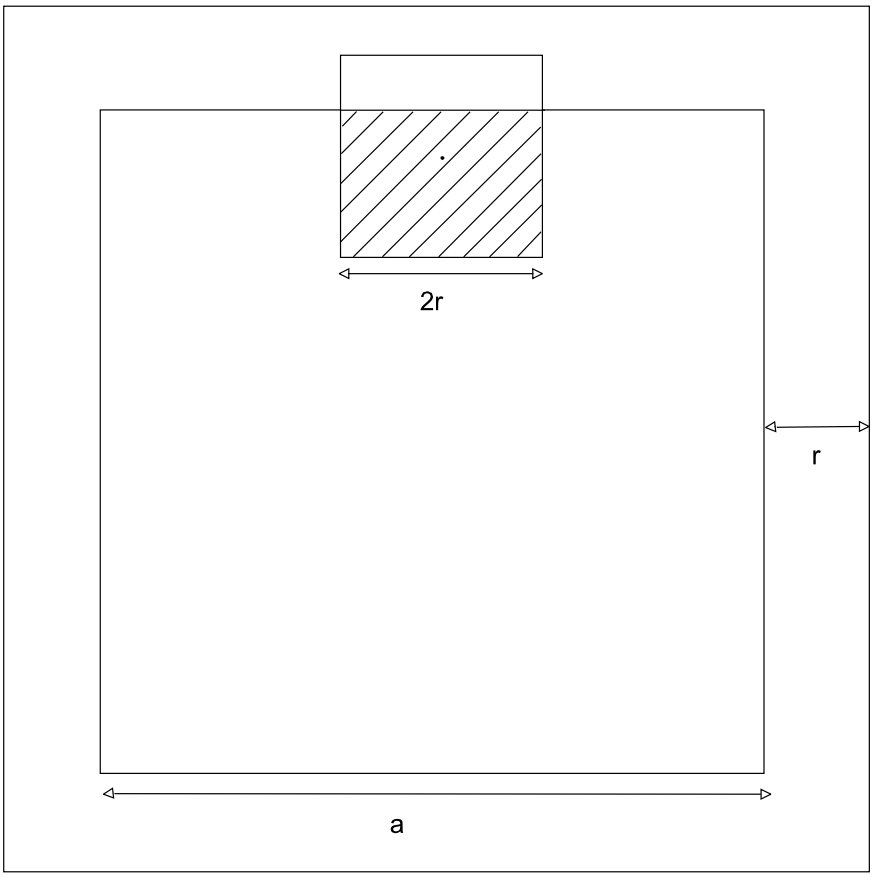

Figure 3

Table 1: The effect on confidentiality of disclosing the variance of normal perturbations; the smaller the table entry, the greater the effect. Each entry in the table is the ratio of the area of confidence region (2) to the expected area of confidence region (4) corresponding to a particular combination of k, n, and α.

k	n	$\alpha = 0.05$	$\alpha = 0.01$
10	2	0.857	0.787
	3	0.927	0.889
	4	0.951	0.926
20	2	0.927	0.889
	3	0.963	0.944
	4	0.976	0.962
30	2	0.951	0.926
	3	0.976	0.962
	4	0.984	0.975

Table 2: Area of (8) when $n = 1$, and estimates of the ratio of expected area of (8) when $n \geq 2$ to the area of (8) when $n = 1$, for different combinations of n, r, and α. Standard errors of estimates are given in parentheses. Results for $n = 2$ are based on 300,000 simulations, and results for $n \geq 3$ are based on 1000 simulations using 10^6 grid points.

n	r	$\alpha = 0.01$	$\alpha = 0.05$
1	0.05	7.775×10^{-3}	7.461×10^{-3}
	0.10	3.110×10^{-2}	2.985×10^{-2}
2	0.05	0.461 (0.000)	0.465 (0.000)
	0.10	0.461 (0.000)	0.466 (0.000)
3	0.05	0.263 (0.005)	0.272 (0.006)
	0.10	0.275 (0.006)	0.269 (0.006)
4	0.05	0.170 (0.004)	0.175 (0.004)
	0.10	0.175 (0.004)	0.176 (0.004)
6	0.05	0.092 (0.002)	0.095 (0.003)
	0.10	0.094 (0.002)	0.093 (0.003)
8	0.05	0.055 (0.002)	0.057 (0.002)
	0.10	0.055 (0.001)	0.056 (0.002)

Table 3: Estimates of the ratio of the expected maximum area of (12) across cases to the expected area of (8) for an arbitrary case, and the coverage probability of (12) for the maximum-area case, for different combinations of k, r, and α, when $n = 2$. Standard errors are reported in parentheses. Results are based on 10,000 simulations.

		Area Ratio		Coverage probability	
k	r	$\alpha = 0.01$	$\alpha = 0.05$	$\alpha = 0.01$	$\alpha = 0.05$
10	0.05	1.030 (0.003)	1.034 (0.003)	0.572 (0.005)	0.559 (0.005)
	0.10	1.025 (0.003)	1.029 (0.003)	0.578 (0.005)	0.554 (0.005)
20	0.05	1.293 (0.003)	1.302 (0.003)	0.684 (0.005)	0.657 (0.005)
	0.10	1.290 (0.003)	1.302 (0.003)	0.675 (0.005)	0.654 (0.005)
30	0.05	1.421 (0.002)	1.440 (0.002)	0.721 (0.005)	0.712 (0.005)
	0.10	1.422 (0.002)	1.439 (0.002)	0.728 (0.004)	0.707 (0.005)

Table 4: Estimates of the coverage probability of (14), for different combinations of σ and α. Standard errors are reported in parentheses. Results are based on 10,000 simulations.

σ	Case	$\alpha = 0.01$	$\alpha = 0.05$
0.05	1	0.990 (0.0010)	0.949 (0.0022)
	2	0.991 (0.0009)	0.952 (0.0021)
	3	0.990 (0.0010)	0.951 (0.0022)
	4	0.991 (0.0009)	0.955 (0.0021)
	5	0.988 (0.0011)	0.952 (0.0021)
0.20	1	0.990 (0.0010)	0.949 (0.0022)
	2	0.989 (0.0010)	0.949 (0.0022)
	3	0.990 (0.0010)	0.950 (0.0022)
	4	0.965 (0.0018)	0.896 (0.0030)
	5	0.968 (0.0018)	0.898 (0.0030)
1.00	1	0.926 (0.0026)	0.840 (0.0037)
	2	0.937 (0.0024)	0.849 (0.0036)
	3	0.928 (0.0026)	0.832 (0.0037)
	4	0.928 (0.0026)	0.834 (0.0037)
	5	0.934 (0.0025)	0.853 (0.0035)

Figure Captions

Figure 1. Depiction of disclosed (left panel) and undisclosed (right panel) case labels for an example in which there are four cases (labeled as 1, 2, 3, 4) and three masked releases.

Figure 2. Depiction of true locations (upper left plot) and typical realizations (remaining plots) when $\sigma = 0.05$, 0.20, and 1.00 for the example having undisclosed case labels described in the text. True locations are at points $(1, -1)$, $(-1, -1)$, $(-1, 1)$, $(1, 1)$, and $(1.25, 1.25)$.

Figure 3. Scenario combining aggregation with perturbation. The outer region is denoted by T, and the square of side a contained in T is denoted by S. The point in the upper portion of S represents a perturbed case location $\mathbf{z} = (x, y)$. The square of side $2r$ centered on $\mathbf{z}$ is denoted by $B(\mathbf{z}, r)$, and the cross-hatched portion of this square, $B(\mathbf{z}, r) \cap S$, has area denoted by $A(\mathbf{z}, r)$. See text for further details.