

REDUCING ESTIMATION BIAS IN ADAPTIVELY CHANGING MONITORING NETWORKS WITH PREFERENTIAL SITE SELECTION.

BY JAMES V. ZIDEK^{1‡} AND GAVIN SHADDICK[†] AND CAROLYN G. TAYLOR^{1‡}

University of British Columbia[‡] and University of Bath[†]

This paper explores the topic of preferential sampling, specifically situations where monitoring sites in environmental networks are preferentially located by the designers. This means the data arising from such networks may not accurately characterize the spatio-temporal field they intend to monitor. Approaches that have been developed to mitigate the effects of preferential sampling in various contexts are reviewed and building on these approaches, a general framework for dealing with the effects of preferential sampling in environmental monitoring is proposed. Strategies for implementation are proposed leading to a method for improving the accuracy of official statistics used to report trends and inform regulatory policy. An essential feature of the method is its capacity to learn the preferential selection process over time and hence to reduce bias in these statistics. Simulation studies suggest dramatic reductions in bias are possible. A case study demonstrates use of the method in assessing the levels of air pollution due to black smoke in the UK over an extended period (1970-1996). In particular, dramatic reductions in the estimates of the number of sites out of compliance are observed.

1. Introduction. This paper addresses the location of monitoring sites within environmental monitoring networks. In many cases, sampling locations may be dependent on the responses themselves or parameters of the underlying environmental process leading to what Diggle et al. [2010] (hereafter referred to as D10) refers to as ‘preferential sampling’. Since measurements from these sites may be critical for informing policy, assessing adherence to standards or for health analyses, the potential effects of such sampling may be of concern. For example urban air pollution monitoring sites provide information that may be used to detect noncompliance with air quality standards [EPA, 2006]. The designer may then locate the sites where air pollution levels are believed to be the highest; although reaching that goal presents its own challenges as shown in Chang et al. [2007]. Reaching this goal would mean the measured concentrations would overestimate the levels of the pollutant in that urban area. That could render these data unsuitable for other purposes, for example in epidemiological studies estimating risks to health.

The focus here is on the calculation of official statistics, where both simplicity and transparency are important. Such statistics traditionally estimate population averages, totals and proportions. Their widespread use and importance means that unbiased estimates are essential and that aggregates of such statistics will also be unbiased. In this paper we propose an approach based on the Horvitz-Thompson (HT) estimator, which has played a key role in the theory of

[‡]The research reported in this article was supported by the Natural Science and Engineering Research Council of Canada.

AMS 2000 subject classifications: Primary 62P12; secondary 62D99

Keywords and phrases: preferential sampling, Horvitz-Thompson estimation, response biased sampling, space-time fields

survey sampling, to producing unbiased estimates of such statistics. The approach builds on the idea of response biased sampling surveyed in Scott and Wild [2011] (hereafter S11), which extends the work in Lawless et al. [1999].

The work was motivated by changes in a large scale air pollution monitoring network in the UK. The network measured Black Smoke (BS), a measure of fine particulate matter, and was in operation for more than fifty years. The very high annual concentrations seen in the early part of its operation led to successful mitigation measures and a dramatic decline in those levels. As a result, the need to monitor BS decreased and the number of sites was reduced. At its peak in the 1960s, there were over 1000 sites but of these only 35 were still operational in the mid-nineties. The sites that were removed from the network had lower concentrations than those that remained and the (small number of) sites that were added to the network over this time had higher levels [Shaddick and Zidek, to appear]. There is therefore clear evidence of preferential sampling over time and thus the decline in reported annual levels of BS was systematically underestimated. This in turn means that exceedances of statutory limits may be overestimated and estimates of health effects of BS may be biased.

A primary aim of this paper is to develop a methodology to adjust annual averages of BS and the proportions of regions in non-compliance with criteria imposed as part of the mitigation strategies for preferential sampling. It is more generally about approaches to correcting for the deleterious effects of preferential sampling on population parameter estimates. In order to address these issues, new theory is developed and assessed.

The remainder of the paper is organized as follows. Section 2 provides a background to previous approaches to preferential sampling in both non-spatial and spatial settings with consideration of issues commonly encountered in environmental modelling. Section 3 provides a superpopulation framework for building a unified approach to dealing with the effects of preferential selection. This is followed in Section 4 by strategies for implementation based on the HT-approach. Section 5 demonstrates the use of the proposed methods using simulation studies which are then applied in Section 6 to the case study of changes in the UK BS monitoring network from 1970 to 1996. In these examples, we show that correcting for bias can substantially reduce estimates of the number of sites, monitored and unmonitored that are out of compliance with regulatory standards. Section 7 discusses our findings and provides some suggestions for alternative approaches to mitigating the effects of preferential sampling in a number of settings.

2. Background. In this section, we review previous approaches that have been proposed to mitigate the effects of preferential sampling. We start with methods that were originally proposed in a non-spatial context but which we contend can contribute to the development of the field and which we relate to a spatial setting. This is followed by a review of methods which have recently been proposed in the field of spatial modelling. The methodology that we propose here, which is unique in this setting in its use of temporal changes to infer the levels of preferential sampling, draws on aspects of many of these approaches and this is discussed at the end of the section.

We note the distinction between the design-based and model-based approaches to spatial design and inference as described by Cicchitelli and Montanari [2012], who argue that the former is appropriate when inference relates to global quantities, such as means or totals, and the latter when ‘constructing a map’, i.e. when performing spatial prediction or interpolation. The approaches that have been used to develop monitoring networks may be classified as: (1) unknown; (2) a combination of networks each developed by an unknown approach [Zidek et al., 2000]; (3) unknown in detail but subject to guidelines; (4) design-based (multi-stage surveys) [EPA, 2009];

(5) model-based [Schumacher and Zidek, 1993]; (6) model-based redesign of networks designed by unknown approaches [Ainslie et al., 2009]. The approaches are diverse or unknown but statistical models can be created to model the results and explore the bias in the outcomes.

Throughout the section, we consider a response, Y , and a set of covariates, or explanatory variables, X , which in a spatial setting would be indexed by their spatial locations u_j . Often interest is in estimating the association between Y and X , represented by β where $f(Y) \sim (\beta_0 + \beta X)$.

Response-biased regression modelling. Here we consider the possibility for bias in estimating relationships between a response Y and a set of covariates, or explanatory variables, X when the sample of data may be subject to preferential sampling. This is referred to as modelling ‘with response-biased samples’ by Scott and Wild [2011] (hereafter S11), who extend the work of Lawless et al. [1999]. This has origins in case-control observational studies where the response Y is a ‘case’ or a ‘control’ with X being observed (and thus available for analysis) for a sample from the population of cases and controls. Models here assume a finite population of possible sample items.

Two approaches for inference are suggested in S11 and we refer to them as follows:

HT: the Horvitz-Thompson (HT) approach (Horvitz and Thompson [1952])

- uses estimating equations designed to construct design-unbiased estimators when finite population elements have unequal probabilities of being selected in a design-based analysis.

CML: the conditional maximum likelihood (CML) approach

- based on the profile likelihood found by maximizing the joint distribution over all possible marginal distributions.

We now describe each of these methods in turn, starting with the HT, and consider how they might adapt to be used in a spatial setting.

Define R as the sampled site indicator such that R_u is 1 or 0 according as site u is selected into the sample or not. Let

$$\pi_u = \pi(y_u, x_u) = P\{R_u = 1 | y_u, x_u\}$$

the selection probability for site u . Here we consider the case of spatial regression where interest will focus on inference on β , the coefficient associated with the explanatory variables in the mean function. Its estimate is found by using the HT-approach and solving the estimating equations

$$(2.1) \quad \sum_u \frac{R_u}{\pi_u} \frac{\partial \log [y_u | x_u, \beta]}{\partial \beta} = 0,$$

assuming $\pi_u > 0, u \in \mathcal{D}$ are known at the sampled sites. When working under the asymptotic paradigm, the selection probabilities are required to be consistently estimable rather than the requirement that they be known.

The second, CMT, approach is based on the profile likelihood found by maximizing the joint distribution over all possible marginal distributions $[X]$ for X . Maximizing the resulting profile yields the estimating equation that characterizes the CML-approach :

$$(2.2) \quad \sum_u R_u \frac{\partial \log [y_u | x_u, \beta, R_u = 1]}{\partial \beta} = 0,$$

which depends on the $\{\pi_u\}$, each being the (conditional) probability that $R_u = 1$.

In their current form both the HT– and CML– approaches have limitations in the setting which we consider here. Firstly, the assumption that the responses Y on which the π_u depend are known seems implausible unlike in case-controlled studies. Their estimation would require a combination of the design– and model– based approaches. Secondly, the failure to incorporate spatial correlation is likely to be a serious limitation. Consistency of the solutions to estimating equations (2.1) and (2.2) means covariances can be estimated by the ‘sandwich estimator’ [Rao et al., 1998] and if samples are sufficiently large this means the assumption of spatial independence may be avoided. However this may be of limited use since many applications such as spatial prediction rely on the estimated covariance and estimates of its parameters.

Cicchitelli and Montanari [2012] propose a design– based approach which specifically addresses the issue of spatial dependence. This is based on the premise that spatial dependence is a result of unmeasured or unrecognized covariates, a view articulated in D10’s discussion section. Sites in close proximity to one another will be influenced by local environmental factors that reflect a spatial pattern and, if not included in a model, these factors will induce spatial pattern. The (measured) covariate matrix, X , is augmented with a matrix of quasi-covariates, Z . For each location, u_j , the quasi-covariate values in Z , $z_{u_j} = (z_{1u_j}, \dots, z_{Ku_j})$, are modelled using thin-plate splines with $z_{u_j} = \tilde{z}_{u_j} \Xi^{-1/2}$, $\tilde{z}_{ku_j} = (||u_j - \kappa_k||)^2 \log (||u_j - \kappa_k||)$, $k = 1, \dots, K$ where the $\{\kappa_k\}$ represent a representative sets of spatial locations (knots) such that Ξ is non-singular. The response is then modelled as a combination of the measured and unmeasured covariates,

$$E_\xi(Y_{u_j}) = \beta_0 + \beta_1 x_{1u_j} + \dots + \gamma_1 z_{1u_j} + \dots + \gamma_K z_{Ku_j}$$

where E_ξ denotes the expected value with respect to the model ξ (in terms of the material introduced in Section 3 this is the superpopulation model’s distribution for Y from which the finite population is drawn). They then invoke the design– based approach and use HT estimators to fit a regression model for Y on (X, Z) . In line with the HT– approach, they do not explicitly model spatial structure in their model assisted, design– based, approach arguing that spatial pattern is provided by the mean function with the augmented covariates. Their position is supported by the well-known duality between first order and second order features of geostatistical models; misspecification of the first order mean function will always lead to bias in the second order variogram.

Spatial prediction. We now consider the model– based approach as used in (D10: Diggle et al. [2010], Gelfand et al. [2012], Pati et al. [2011]) all of which specifically consider spatial modelling. D10, which includes a bibliography of earlier work on preferential sampling unrelated to this paper, characterizes the effect of preferential sampling within a model– based geostatistical framework in which site locations are informative for inference. D10 assumes a latent, unobservable Gaussian field S over a geographical continuum (domain) \mathcal{D} . The sites u are selected at random in accordance with an inhomogeneous Poisson spatial process with intensity function $\lambda_u = \exp\{\alpha + \beta S_u\}$, $u \in \mathcal{D}$. The measurable response Y is also modelled as dependent on S . Both the response measurements, y^{obs} , and selected sites U yield information about the underlying model parameters, including both those in the spatial mean as well as spatial covariance matrix for S . A marginal likelihood function is obtained by marginalizing out S although S , U and Y are correlated in their joint spatial distribution. A key to the success of the model in D10 is knowledge of process and data, both present and future, in order to characterize the sampling selection process, an assumption questioned by Dawid [2010]. Given that knowledge, the effects

of preferential sampling on variogram estimates, spatial predictions and associated biases can be assessed. Pati et al. [2011] extend that approach, again relying on latent variables in a point process approach. Gelfand et al. [2012] suggests an alternative approach based on knowledge of the underlying process being monitored and the factors that drive that process.

The temporal dimension. The approaches described above do not include changes to networks due to preferential sampling over time, a key feature of spatial sampling addressed in this paper. Examples of this include the redesign of Vancouver’s air quality monitoring network [Ainslie et al., 2009] and the case study in Section 6 which considers changes in a long-term air pollution monitoring network in the UK. Le and Zidek [2006] demonstrate the use of such adaptation as the network for monitoring ground level ozone concentrations that has been steadily augmented over the last few decades as the adverse health effects of ozone have been recognized. In this example, it seems plausible that the addition of sites has been performed preferentially to ensure high levels of ozone are detected. The case study in this paper considers the monitoring of black smoke (BS) for which the number of monitoring sites declined from ca. 1000 in the early 1960s to ca. 100 at the turn of the last century as levels of BS declined due to improvements in the management of air quality. Shaddick and Zidek [to appear] demonstrate that the reduction in the network was done preferentially by removing sampling sites with generally lower concentrations relative to the decline in overall levels of BS, leading to overestimates of BS concentrations (as can be seen in Section 6). We propose that by considering the stochastic process of selection over time, it is possible to model the extent of preferential sampling and estimate its deleterious effects on published estimates of environmental fields and the effects on items of inferential interest. This emphasis leads us to build on the HT–approach and the use of estimating equations following the celebrated work in this area of David Binder [Binder, 1983].

Unlike the earlier methods previously described, this approach incorporates time in the model. We draw on the previous methods and specifically those based on the H–T estimator, and assume that the sampling probabilities are uncertain and model the selection process as stochastic, depending sequentially on the responses over time. The result is an approach that is a combination of the design– and model–based approaches. The $\{\pi\}$ are then learned over time as the results of monitoring accumulate.

Populations of potential sampling sites can often be taken, as in this paper, to be a finite set of possible locations. In that sense, this approach diverges from D10 and its successors in the model–based category by assuming a finite population of N sampling (i.e. monitoring) sites, u_j , $j = 1, \dots, N$. Note that u_j could represent just the label j or more commonly the geographic coordinates of the site, depending on the context. This is not to say that the assumption of a continuous domain for site selection is unreasonable in all cases but practical and administrative considerations will often restrict S to be a finite dimensional vector–valued process over a discrete domain \mathcal{D} . In D10 S is also discretized but this is in order to approximate the marginal likelihood and is done by replacing the continuous \mathcal{D} by a fairly dense lattice. Sampling points then have to be mapped onto their nearest lattice point neighbours.

As with the successors to D10 cited above, we allow covariates to be incorporated. Site selection may well depend on them, for example the distance of a site from a major roadway or an ‘urban–rural’ classification (as used in the case study in Section 6). In addition, interest may well lie in the significance of the effects of such covariates or design variables on the measured responses.

3. A general framework. Section 2 presents a number of paradigms in which to study the issue of preferential spatial sampling within the design- or model-based frameworks. The latter includes the Bayesian approach although it was not explicitly mentioned. The concept of a superpopulation [Sarndäl et al., 2003] provides a framework for unifying paradigms for inference and it is within this framework that we develop the proposed approach. We combine design- and model-based approaches, allowing us to estimate the unknown selection probabilities needed for the HT estimator which are used to compensate for the selection bias introduced in sampling from the finite population of interest.

We suppose a discrete geographical domain \mathcal{D} contains point referenced sites $u_j, j = 1, \dots, N$. Let T denote the present time and although the spatial locations do not have a natural ordering it is convenient to use the vector notation $\mathbf{Y}_t : 1 \times N$ to represent the sequence of responses at those sites at times $t = 1, \dots, T$. These sites, which need not be on a lattice, represent a finite population of potential locations at which to site monitors that repeatedly measure at regular times, a random space-time field. Further let \mathbf{Y} denote the $T \times N$ matrix comprised of those row response vectors. Similarly at time t , let \mathbf{X}_t denote a matrix of covariates or explanatory factors, hereafter referred to as ‘covariates’ for simplicity. Then \mathbf{X} denotes the corresponding three dimensional covariate array.

We now propose a framework which contains three major components in the process of using data from monitoring networks which for simplicity in subsequent descriptions we characterise as (i) Nature, (ii) the Preferential Sampler and (iii) the Statistician;

- (i) Nature governs the process-model; a joint distribution for \mathbf{X} and \mathbf{Y} that generates realizations \mathbf{x} and \mathbf{y} over the time period ending at the present time, T . These are regarded as drawn from an infinite population of possibilities called the Superpopulation that is indexed by the sites in \mathcal{D} . In some contexts this would be the relevant population and the parameters of the process-model the objects of inferential interest. However here we consider the case, commonly encountered in official statistics, where there is a finite population of possible site locations and thus values of \mathbf{x} and \mathbf{y} which we refer to as the Population.
- (ii) The Preferential Sampler runs the measurement-model; the process of chooses sites in \mathcal{D} at which the associated elements of \mathbf{x} and \mathbf{y} are observed. The sampling design, i.e. the selection probabilities for sites to be included in the sample at each time $t = 1, \dots, T$ may depend on the elements of \mathbf{x} and/or \mathbf{y} for previous times. The resulting sample can be interpreted as from either the Population or Superpopulation depending on the goal of the monitoring program.
- (iii) The Statistician, working within a design-based framework and using only the sample and knowledge of the sites at which the sample was collected, infers the uncertain selection probabilities used by the Preferential Sampler and through this adjusts inferences about the Population (or Superpopulation as appropriate) to compensate for the bias induced by preferential sampling.

The mechanisms of the preferential sampling process and associated selection probabilities can be complex and are often not well understood or may be unknown. They may non-anthropogenic or anthropogenic. Examples of the former might include a natural event, such as a forest fire, making some locations inaccessible (and thus their selection probabilities zero). Although responses continue to be generated they could not be sampled. For the latter, there may be new guidelines on where sites may be located or new developments or changes of land ownership may mean sites can no longer be located at the same places. These uncertain mechanisms that lead

to the sample make the selection probabilities uncertain as well meaning that they need to be inferred. The Superpopulation-Population-Sample paradigm provides a framework where that becomes feasible.

We now formalize the ideas above in a general theory for response generation (\mathbf{Y}) and site selection (\mathbf{R} from Section 2), one that allows flexibility in the choice of modelling, inferential and selection paradigms. We first develop a theory based on the conditional distribution of \mathbf{Y} given $\mathbf{X} = \mathbf{x}$ and model parameters θ . These model parameters would include all those that characterize the joint distribution and could include, depending on the context, regression coefficients β (as seen in Equation 2.1) and autocorrelation and spatial covariance parameters. Using the notation from D10 we denote the conditional distribution of \mathbf{Y} , which may be characterized by its probability density or cumulative distribution function, by

$$(3.3) \quad [\mathbf{y}|\mathbf{x}, \theta].$$

If within a specified time period, all responses and covariates were observed for every spatial site in the Population, we could proceed in the usual way to make inferences about θ . In particular, given $\mathbf{Y} = \mathbf{y}$ and $\mathbf{X} = \mathbf{x}$ the conditional likelihood function would be given by Equation (3.3). The Superpopulation's maximum likelihood estimator (MLE) of θ , denoted by $\hat{\theta} = \theta(\mathbf{y}, \mathbf{x})$ would estimate θ , including temporal as well as spatial correlation model parameters together with coefficients in the regression model relating \mathbf{Y} and \mathbf{X} . Alternatively if only a random sample of sites were selected from the finite population of sites, and their associated response–covariate pairs recorded, an estimate $\hat{\hat{\theta}}$ of $\hat{\theta}$ could be computed and considered to be an estimate of θ .

In contexts where official statistics are collected and published or regulatory policy is administered, estimates for specific times are commonly required. In such cases the Population at the present time T consists of just the responses $\mathbf{Y}_T = \mathbf{y}_T$ generated by the marginal distribution

$$(3.4) \quad [\mathbf{Y}_T | \mathbf{X}_T = \mathbf{x}_T, \theta_T],$$

where we have assumed \mathbf{Y}_T depends only on the covariates at time T . Hereafter for simplicity ‘Population’ will refer generically to this time dependent population. Following a standard approach in survey sampling theory [Sarndäl et al., 2003], we would define $\hat{\theta}_T$ to be the matrix of parameters of the Population, obtained by maximizing the marginal likelihood in Equation 3.4 and take it to be the object of inference, although it may be also be viewed as representating θ_T . (For a discussion of this issue see Pfeffermann [1993].) Therefore two legitimate objects of inferential inference present themselves, θ_T and $\hat{\theta}_T$. In either case, as in D10, we are concerned with the effects of preferential sampling on the estimates derived from the sample of $\hat{\theta}_T$ and in turn θ_T , depending on the inferential objective.

To formalize these ideas, we express the Superpopulation log-likelihood estimating equation for the MLE $\hat{\theta}_T$ as

$$(3.5) \quad \nabla_{\theta_T} \log [\mathbf{y}_T|\mathbf{x}_T, \theta_T] = 0.$$

The measurement–model is more complex since the process for selecting the sites at time T , on which inference is to be based at that time, may depend on responses at previous times. To model the selection process we use notation introduced in Section 2. Thus we let \mathbf{R} denote the $T \times N$ matrix of indicator random variables whose t^{th} -row of \mathbf{R}_t consists entirely of zeros except for ones in the columns corresponding to the sites selected for inclusion at time t .

Let \mathbf{y}_r and \mathbf{x}_r denote the observed values of \mathbf{Y} and \mathbf{X} at the design points selected adaptively over time. In other words if $\mathbf{r} = (r_{tj})$, then

$$\mathbf{y}_r = \{y_{tu_j} : t, j \text{ for which } r_{tj} = 1\}$$

and so on. We model the distribution of \mathbf{R}_t as stochastically dependent only on $\mathbf{Y}_{1:(t-1)}$ and $\mathbf{X}_{1:(t-1)}$, where we use the general notation, $a_{r:s} = (a_r, \dots, a_s)$ for $r \leq s$ and the null vector if $r > s$ for any object a . That dependence could reflect the effect of a latent process as in D10. Then $\pi_{tu} = P(R_{tu} = 1 | \mathbf{y}_{1:(t-1)}, \mathbf{x}_{1:(t-1)}, \eta)$, $t = 1, \dots, T$, where η denotes the matrix of parameters for the measurement-model. Our assumptions imply that the conditional preferential sampling distribution of \mathbf{R} is given by

$$(3.6) \quad [\mathbf{r} | \mathbf{y}, \mathbf{x}, \eta] = \prod_{t=1}^T [\mathbf{r}_t | \mathbf{y}_{\mathbf{r}_{1:(t-1)}}, \mathbf{x}_{\mathbf{r}_{1:(t-1)}}, \mathbf{r}_{1:(t-1)}, \eta].$$

Combining Equations (3.4) and (3.6), yields for inference at time time T , the joint conditional likelihood

$$(3.7) \quad L(\eta, \theta_T) \doteq [\mathbf{y}_T | \mathbf{x}_T, \theta_T] \prod_{t=1}^T [\mathbf{r}_t | \mathbf{y}_{\mathbf{r}_{1:(t-1)}}, \mathbf{x}_{\mathbf{r}_{1:(t-1)}}, \mathbf{r}_{1:(t-1)}, \eta].$$

We assume that η does not contain elements of θ_T and so the population parameter matrix remains as that defined in Equation 3.5. The likelihood here suggests an approach for fitting the site selection probabilities once the sample is obtained: impute the non-sampled values and estimate η from the likelihood inferred from the resulting combination of actual and imputed data. This is the approach used in later sections with a logistic regression approach. Note that while the preferential sampling scheme represented in Equation 3.7 is ancillary for the purpose of estimating of θ_T , within the design-based framework below for inferring the population parameter θ , it is very relevant and in fact it lies at the heart of the HT-approach used there.

Our general framework can be extended to include a Bayesian approach by incorporating a prior joint distribution for η and θ . This topic is left for future work.

General implementation strategies. We now describe general strategies for implementing our general framework using the HT-approach in Section 2, leaving Section 4 for specific techniques. We demonstrate how the general framework might work and provide and link to what follows in the next section where the HT-approach is developed. More specifically we show that the framework can be used even when we cannot obtain estimating equations resembling those in (2.1) when inter-site dependence is present.

The population parameter matrix associated with that of the superpopulation process, $\theta = (\theta_1, \dots, \theta_T)'$, commonly has the form

$$(3.8) \quad \hat{\theta}_t = \mathbf{H} \left\{ N^{-1} \sum_j (\mathbf{h}_1[\mathbf{y}_{tu_j}, \mathbf{x}_{tu_j}], \dots, \mathbf{h}_q[\mathbf{y}_{tu_j}, \mathbf{x}_{tu_j}]) \right\}.$$

for known functions \mathbf{H} and $\mathbf{h}_1, \dots, \mathbf{h}_q$.

Then if the $\{\pi_{tu_j}\}$ are known or estimated, $\hat{\theta}$ can be estimated by

$$(3.9) \quad \hat{\theta}_t = \mathbf{H} \left\{ \sum_u \frac{r_{tu}}{N\pi_{tu}} (\mathbf{h}_1[\mathbf{y}_{tu}, \mathbf{x}_{tu}], \dots, \mathbf{h}_q[\mathbf{y}_{tu}, \mathbf{x}_{tu}]) \right\}.$$

Justification for this choice comes from the unbiasedness of $\sum_u (r_{tu}/N\pi_{tu})\mathbf{h}_l[\mathbf{y}_{tu}, \mathbf{x}_{tu}]$, $l = 1, \dots, q$ as estimates of their corresponding population averages.

We illustrate this approach in two specific cases: (1) a regression model where the interest is in the association between \mathbf{Y} and \mathbf{X} , a relationship that may evolve over time; (2) the estimation of population means in the presence of spatial dependence.

Case 1. Regression of \mathbf{Y} on \mathbf{X} : Here

$$H(\mathbf{a}, \mathbf{b}) = \begin{pmatrix} a_1/b_1 \\ \vdots \\ a_T/b_T \end{pmatrix}$$

for T dimensional vectors \mathbf{a} and \mathbf{b} with $h_1[y_{tu}, x_{tu}] = y_{tu}x_{tu}$ and $h_2[y_{tu}, x_{tu}] = x_{tu}^2$.

Case 2. Spatially dependent Gaussian fields: Conditional on the mean and covariance structure, the Superpopulation model is given by a matrix–Normal distribution. That is

$$\mathbf{Y} \sim N_{T \times N}(\boldsymbol{\mu}, \boldsymbol{\Omega} \otimes \boldsymbol{\Lambda})$$

where $\boldsymbol{\Omega}^{T \times T}$ as well as $\boldsymbol{\Lambda}^{N \times N}$ are positive definite matrices, $E(\mathbf{Y}_{tu}) \equiv \nu_t$, $t = 1, \dots, T$ for all u , and $\boldsymbol{\mu} = \boldsymbol{\nu} \otimes \mathbf{1}$ with $\boldsymbol{\nu} = (\nu_1, \dots, \nu_T)'$, $\mathbf{1}^{1 \times N} = (1, \dots, 1)$. Assuming $\boldsymbol{\Omega}^{-1} = \boldsymbol{\tau}\boldsymbol{\tau}'$ is known, the population parameters, $\{\hat{\nu}_t\}$ are easily found to be

$$(3.10) \quad \hat{\nu}_t = \sum_j \frac{\mathbf{y}_t \boldsymbol{\tau}^j \boldsymbol{\tau}^{j'} \mathbf{1}'}{\sum_j (\mathbf{1} \boldsymbol{\tau}^j)^2} = \frac{\mathbf{y}_t \boldsymbol{\Omega}^{-1} \mathbf{1}'}{\mathbf{1} \boldsymbol{\Omega}^{-1} \mathbf{1}'}$$

where $\boldsymbol{\tau} = (\tau^1, \dots, \tau^N)$. This population parameter has the form given in Equation (3.8) and so can be adjusted using the HT–approach. Note that the profile likelihood for the covariances $\boldsymbol{\Omega}$ and $\boldsymbol{\Lambda}$ involves a quadratic form in \mathbf{y} , and it yields estimating equations for the population level MLEs which can be used to estimate them.

This case is more general than it may seem at first glance. The Gaussian likelihood in this example can be treated as a quasi likelihood leading to the GEE approach [Liang and Zeger, 1986] when \mathbf{Y} is not Gaussian. The so–called working covariance between columns can be taken as independent in that case with asymptotic justification providing that the assumption of equal means across sites holds. Alternatively we can model spatial patterns via the covariates, \mathbf{X} , and then assume no spatial structure in the covariance as in Cicchitelli and Montanari [2012] (see Section 2).

Another approach to implementing the general framework is also available when considering regression where an estimating equation of the form seen in (2.1). For this case we suppose $\theta_t = (\boldsymbol{\beta}_t, \boldsymbol{\Psi})$ where $\boldsymbol{\beta}_t$ is the vector of regression parameters. This plays the same role as it did in Equation 2.1, that of parametrizing first order effects embraced by the process mean function, while $\boldsymbol{\Psi}$ represents the parameters of the spatial dependence model, i.e. the covariance in the case of a Gaussian field. We denote $\boldsymbol{\Psi}_0^{\text{known}}$ as the case when there is spatial independence, a diagonal matrix in the case of the Gaussian field. Note that Section 2 provides a discussion of the assumption of independence.

From Equation 3.5, we now get the superpopulation maximum likelihood estimating equation

for inference about the population at time T :

$$\begin{aligned}
 (3.11) \quad & \nabla_{\beta_T} \log [\mathbf{y}_T | \mathbf{y}_T, \mathbf{x}_T, \beta_T, \Psi_0] \\
 &= \sum_u \nabla_{\beta_T} \log [y_{Tu} | y_{Tu}, x_{Tu}, \beta_T, \Psi_0] \\
 &= 0.
 \end{aligned}$$

In this way the population parameters are defined.

Following S11, we now have two approaches for finding the sample-based MLE. The first is provided by the HT-approach in Equation 2.1, which yields the following estimating equation

$$(3.12) \quad \sum_u \frac{r_{Tu}}{\pi_{tu}} \nabla_{\beta} \log [y_{Tu} | x_{Tu}, \beta_T, \Psi_0] = 0.$$

The CML-approach for estimating β would rely on the following estimating equation from Equation 2.2 where

$$(3.13) \quad \sum_u \nabla_{\beta_T} \log [y_{Tu} | x_{Tu}, \beta_T, \Psi_0, R_{tu} = 1] = 0.$$

If responses can be assumed to be temporally independent then this simplifies to be of the same form as (3.12)

A third approach to implementing our framework involves generalizing the maximum likelihood estimating equation to an general estimating equation as described in Godambe and Thompson [1986]. In the case of temporally independent fields, for the superpopulation case, this becomes

$$(3.14) \quad \sum_u \phi_{Tu}(y_{Tu}, x_{Tu}, \beta_T) = 0$$

for some known functions $\{\phi_{Tu}\}$. Note that under regularity conditions this gradient has a conditional expectation, given the superpopulation parameters, equal to zero; a property referred to an ‘unbiasedness’. In fact, Equation (3.14) can be used to define an estimator for any choice of kernel ϕ provided it is unbiased. In this way, Binder and Zdenek [1994] formulated a general approach for complex sample surveys based on estimating equations.

4. Implementing the HT-approach. This paper proposes the HT-, or ‘inverse probability weighting’, approach to compensate for the bias introduced by preferential sampling. The approach can be implemented in a variety of ways depending on the context and inferential paradigm [Kloog et al., 2012]. A question that arises in all cases is the role of spatial dependence and where or not it needs to be explicitly acknowledged in the chosen model. The approaches described in Section 2 suggest the answer to this question is not clear cut and the answer will depend on various factors including data that are available and the inferential objectives. Even where spatial dependence should be incorporated in the superpopulation model, the specific application may prohibit it. In the example of Case 2 (under the assumption that the model is correct), then the unequally weighted population mean in Equation (3.10) which incorporates

spatial correlation should ideally be estimated. In practice however equally weighted average is commonly used (and is the one that is available for the case study presented in Section 6). In such cases, the HT-approach must be applied to the available estimate rather than an idealized one. This section suggests an approach that covers all situations that fall within the general framework of Section 3 whereas Sections 5 and Section 6 focus on the equally weighted average in order to provide a solution to the problem most likely found in practice.

In practice the process which selects the monitoring sites is non-random and generally not known. The selection probabilities cannot therefore be characterized as they are in multi-stage survey sampling, for example. However having a time series of samples of sites from the finite population of N sites enables us to model the selection process and estimate those probabilities which, in order to account for their uncertainty, are treated as random. We now describe a general logistic regression approach for that purpose.

Assume that at time t , the sample of sites S_t among the population of N sites is selected by a PPS (probability proportional to size) sample survey design, $u \in S_t$ being included with probability π_{tu} . That probability is assumed to depend on all responses, both observed and those unmeasured over the time period $1 : (t-1)$ (the latter being treated as latent variables, analogous to the S s in D10). Thus in terms of the measured and unmeasured responses, Y and vector of binary indicators of selected/rejected sites R , the conditional distribution of the probability of selection is

$$(4.15) \quad \begin{aligned} \text{logit}[\pi_{tu}] &= \text{logit}[P(R_{tu} = 1 \mid \mathbf{y}_{1:(t-1)}, \mathbf{r}_{1:(t-1)})] \\ &= G(\mathbf{y}_{1:(t-1)}, \mathbf{r}_{1:(t-1)}) \end{aligned}$$

for some function $0 \leq G \leq 1$. That function is our analogue of the preferential sampling intensity in Assumption 2 in D10.

Under the assumption of the super-population model there will be a predictive probability distribution for the unmeasured responses. Values for these might be obtained using, for example geostatistical methods, which under repeated imputation will allow $k = 1, \dots, K^*$ replicate datasets. Each replicate enables us to fit G through logistic regression to get \hat{G}^k , $k = 1, \dots, K^*$ and in turn $\hat{\pi}_{tu}^k$, $k = 1, \dots, K^*$. This approach is analogous to Equation (9) in D10, with R playing the role of X .

From these replicates, multiple values of the HT estimator can be obtained, which allows an adjusted point estimate of the population average to be found together with error bands around the point estimate to reflect the associated uncertainty.

The exact way in which these are computed would depend on the scheme by which the network was adapted as illustrated in Section 5. We now discuss the case of a network that expands monotonically over time, referred to as an ascending staircase design. This forms the basis of one of the simulation studies in Section 5 together with one based on a shrinking network.

4.1. *Expanding networks..* Consider the case of an ascending staircase design so that the network expands monotonically over time [Le and Zidek, 2006] i.e. $S_{(t-1)} \subseteq S_t$, $t = 1, \dots, T$ where S_0 is the null set and $S_T \subseteq \mathcal{D}$. After a site, u , has entered the network it remains. Using an * to denote unconditional probabilities, we assume that initially sites in $S_1 \subset \mathcal{D}$ are selected without replacement so that the selection probabilities are $\pi_{1u}^* = \pi_{1u} = |S_1| / N$ where in general $|A|$ is the number of elements in a set A . The HT estimator for the population mean can now

be calculated as

$$(4.16) \quad \hat{\mu}_1 = \sum_{u \in S_1} \frac{y_{1u}}{N\pi_{1u}^*}$$

At time 2 additional sites $S_2 \setminus S_1$ must be selected from $\mathcal{D} \setminus S_1$ and this is assumed to be done with probabilities proportional to size at time 1, i.e. based on the $\{y_{1u}\}$ for these sites. At time 2, the HT–approach sees a single sample of sites S_2 . Inclusion in S_2 means a site was either selected at time 1, in which case it is certain to be in S_2 or it was selected at time $t = 2$ for the first time. Hence overall for all sites in $u \in S_2$

$$(4.17) \quad \pi_{2u}^* = \pi_{1u}^* + (1 - \pi_{1u}^*)\pi_{2u}.$$

However at time 2, the π 's are unknown unlike those at time 1 due to the unknown responses at time 1 on which the preferential sampling was based and the unknown monotone function of these responses implicitly used by the preferential sampler. The unknown responses can be multiply imputed by standard geostatistical (or other) methods to get $\{\hat{y}_{1u}^k, \mathcal{D} \setminus S_1\}$ on replicate $k = 1, \dots, K$ so that in effect, we have a complete set of responses over \mathcal{D} .

Logistic regression can then be used to estimate the probabilities of selection by fitting

$$(4.18) \quad \text{logit}(\hat{\pi}_{2u}^k) = \alpha_2 + \beta_2[\hat{y}_{1u}^k - \bar{y}_1^k],$$

to the N binary select-reject indicators for all sites $u \in \mathcal{D}$ where \hat{y} represents either observed or imputed values at time 1 as appropriate. Replicate HT estimators at time 2 are obtained by multiply imputing the unknown responses and generating multiple HT estimates:

$$(4.19) \quad \hat{\mu}_2^k = \sum_{u \in S_2} \frac{y_{1u}}{N\hat{\pi}_{2u}^{*k}}.$$

Their average yields the adjusted point estimator at time 2 and their empirical quantiles provide a means of estimating (95%) error bands for the true population mean.

We can proceed in a similar fashion at time 3. Here the required imputation of the unobserved y_{2u} is more complicated due to the preferential sampling effect. To do this we adapt a key idea in Pati et al. [2011] and use logit transformations of the estimated selection probabilities as covariates that represent the preferential selection effect. These can then be incorporated in the spatial trend (mean field) model for a Gaussian random field superpopulation model. Since at time 1 we assume no preferential selection, we let the required covariates be $z_{1u} = \text{logit}(\pi_{1u})$, $u \in \mathcal{D} \setminus S_1$. Then at time $t = 2$ we get

$$(4.20) \quad z_{2u} = \text{logit}(\bar{\pi}_{2u}) \quad u \in \mathcal{D} \setminus S_1$$

by averaging the replicate values in Equation (4.18). For completeness take $z_{2u} = z_{1u}$, $u \in S_1$. Note that the z 's correspond to what Pati et al. [2011] calls “plug-in” estimates.

We assume a Gaussian random field superpopulation model with a Matérn covariance matrix and spatial mean field

$$(4.21) \quad E[Y_{2u}] = \mu_{2u} + \vartheta z_{2u}, \quad u \in \mathcal{D}$$

where ϑ is an unknown regression parameter. This model is fit, including any unspecified parameters in μ_{2u} , using for example geostatistical methods and by multiply imputing the unobserved

values of the unobserved $\{y_{2u}\}$. Proceeding recursively in this way leads to a $K \times t$ matrix of replicates $\hat{\pi}_{tu}^{*k}$ for each u at time t . That in turn yields replicates of the HT estimates for the population mean at time t .

Refinements of this approach would be possible, including the addition of a term in Equation (4.20) to incorporate spatial structure. This would be appropriate if it were known that designers took spatial considerations into account, in coming to their new-site admission decisions and in calculating their summary statistics. Model selection presents another challenge for Equation (4.15). Selecting appropriate predictors from the class of all possible metrics that could be computed from previous exposure data will be challenging. Formal model selection approaches will generally be impractical, necessitating reliance on some context specific knowledge to help reduce the class of possibilities.

4.2. Reducing networks. The case in which a network is monotonically decreasing is somewhat simpler as the need for imputation of unmeasured responses is eliminated. Starting from a set of N sites, S_1 at time 1, then at each time t , the set of sites that remain in the network will be $S_t \subset S_{t-1}$. At each time point, the selection probabilities are obtained from logistic regression and the HT estimates then constructed based on Equation (3.9).

The efficacy of the cases of decreasing and expanding networks are explored using simulation in Section 5 with the former also being the basis of the case study presented in Section 6.

5. Simulation study. This section describes simulation studies that explore the approaches described in Section 4. Using the terminology introduced in Section 3 we generate data for the underlying environmental field (Nature) following the structure used in Gelfand et al. [2012]. Given this field, we simulate the role of the Preferential Sampler who at each time t selects sites for inclusion at time $t + 1$ with selection probabilities proportional to the magnitude of the measurements (PPS). The Statistician, having only the measurements at the selected monitoring sites at time t , adopts a superpopulation model for all unmeasured response and, knowing of the use of PPS, fits a logistic regression model to the binary site selection process for time $t + 1$ to estimate the site selection probabilities. The HT estimator is then used to adjust the annual average estimates for time $t + 1$ for the effects of preferential sampling.

The underlying environmental field: Nature. To generate emissions over space and time we consider two cases from Gelfand et al. [2012]. The first of these considers emissions arising from a point source (of pollution) while the second considers pollution from three cities. Central to the approach advocated in Gelfand et al. [2012] is that measured concentrations are based on emissions, which in the case of the three cities are represented by population densities. We generate data over 25 years for a finite population of 1000 sites.

For the point source example, emissions are represented by a kernel x_u given by

$$(5.22) \quad x_u = \exp(-1.8 \|u - q\|), \quad u \in \mathcal{D}$$

with $q = (0.25, 0.75)$. The maximum and minimum values of x are 1 and 0.16, respectively. The second scenario involves three cities with centers located at $c_1 = (0.75, 0.75)$, $c_2 = (0.25, 0.25)$, $c_3 = (0.75, 0.25)$ and their population densities given by

$$(5.23) \quad p_u = \exp(-5 \min_i \|u - c_i\|), \quad u \in \mathcal{D}.$$

The maximum and minimum values of p are 1 and 0.019, respectively.

TABLE 1
Spatial parameters used in the simulation studies.

Model	Nugget	Sill	Range	Smoothness
M1	0	0.0079	0.5	0.5
M2	0	0.00013	0.5	0.5

The following adapts this approach to the cases considered here by incorporating time. At each time t , we assume a linear relationship for the mean concentrations and emissions. For the two cases (point source, M_1 , and multi-city, M_2)

$$\mathbf{M}_1: \mu_{1tu} = \phi_1 x_{tu}$$

$$\mathbf{M}_2: \mu_{2tu} = \phi_2 p_{tu}$$

The emission levels are assumed to decline over time, $t = 1, \dots, T, u \in \mathcal{D}$:

$$(5.24) \quad x_{tu} = x_u - \gamma_{1u}(t - 1)$$

$$(5.25) \quad p_{tu} = p_u - \gamma_{2u}(t - 1),$$

where decay parameters, γ_{iu} , are site specific; $\gamma_{iu} = [a_i \mu_{iu} + b_i] \mu_{iu}$ for the two models M_i ($a_1 = 0.009391$, $b_1 = 0.001216$, $a_2 = 0.008156$, $b_2 = 0.003686$ - these values were chosen to ensure that after simulated Gaussian residuals are added to the spatial mean, the simulated responses will be non-negative (with high probability) over all sites and times). Using this approach, the large mean values are greatly reduced (50% for the largest mean) unlike the small ones (10% for the smallest mean). In the following simulations, the relationship between emissions and concentrations is set to $\phi = 2$.

We select an irregular grid of 1,000 points from a regular grid of 10,000 points to represent the population of possible sampling sites. For each time t , the pollution field is simulated as a Gaussian field with means μ_{jt} as described above and a fixed Matérn spatial correlation structure whose sill parameter differs for the two mean models (again to ensure a high probability of nonnegative simulated values)

The correlation between the process responses for sites separated by a distance d is given by

$$\rho_\theta(d) = \frac{2}{2^\kappa \Gamma(\kappa)} \left(\frac{d\sqrt{2\kappa}}{\omega} \right)^\kappa K_\kappa \left(\frac{d\sqrt{2\kappa}}{\omega} \right)$$

for both models M_i , $i = 1, 2$. Here κ is the smoothness parameter, which in this case is $1/2$ to yield the exponential spatial correlation function. The ω determines the range of the model. The correlation would be multiplied by the sill to get the spatial covariance function. The values of these parameters are specified in Table 1. Thus at each time point a random vector of responses $Y_t^{1 \times 1000} \sim N_{1000}(\mu_t, \Sigma)$ is generated, where Σ is determined by the member of the Matérn family. The result: a matrix of simulated emission levels of dimension 1000×25 which constitute the finite population to be studied in this section.

Selection procedures: The Preferential Sampler. We consider both of the cases described in the previous section; (i) descending (network reducing in size) staircase and (ii) ascending (network increasing) staircase of adaptive network design.

TABLE 2

Parameters used to characterize mild and strong preferential sampling for each of the two emission scenarios considered in this simulation study.

Scenario	Mildly Preferential		Strongly Preferential	
	α	β	α	β
Point source	0.32	0.10	0.32	2
Multi-city	0.038	0.010	0.038	0.5

(i) Shrinking adaptive network design

1. For time $t = 1$, let $S_1 = \mathcal{D}$ be the entire population of N sites.
2. For each successive time $t = 2, \dots, T$ draw a sample $S_t \subset S_{t-1}$ of size $m_t = |S_t| = 25$ with sampling probabilities proportional to size, $\alpha + \beta y_{tu}$, $u \in S_{t-1}$. For each of the two emission scenarios, two cases are considered, the first being that sampling is mildly preferential and the second strongly preferential. The respective selection parameters are given in Table 2.
3. Repeat Step 2 1000 times to generate point estimates and 95% error bars for the annual estimates the selections produce.

(ii) Expanding adaptive network design

1. time 1: Draw a simple random sample of $m_1 = 50$ sites S_1 without replacement from \mathcal{D} of $N = 1000$ sites. Lack of knowledge at that stage makes this selection model plausible.
2. time $t = 2, \dots, 25$: Draw a sample $S_t \supset S_{t-1}$ by adding an additional $m_t = 10$ sites selected (from the remaining unselected sites from the original set of 1000 sites) with probability proportional to size, again $\alpha + \beta y_{tu}$, $u \notin S_{t-1}$.
3. Repeat Step 2 1000 times.

Correcting for preferential sampling: The Statistician. This is done assuming a Gaussian random field (GRF) superpopulation random response, that sampling probabilities are proportional to size and the multivariate response vectors not autocorrelated. For clarity of exposition, we drop the subscript j denoting scenario model M_j . The following are the steps needed for the two cases.

(i) Shrinking adaptive network design

1. time 1: $S_1 = \mathcal{D}$.
2. time $t = 2, \dots, T$: Use logistic regression as in Equation (4.18) to estimate coefficients of the selection model (α and β) and hence get an estimate of the conditional selection probability $\hat{\pi}_{tu}$ for that time. All relevant data are available so no imputation is needed.
3. time $t = 1, \dots, T$: Compute the unconditional site selection probabilities $\hat{\pi}_{tu}^*$ and the HT estimate of the annual mean. Note: $u \in S_t$ implies $u \in S_{t'}$, $t' \leq t$ so under the assumption of no autocorrelation

$$\hat{\pi}_{tu}^* = \prod_{t'=1}^t \hat{\pi}_{t'u}.$$

4. Repeat steps 2 and 3 to compute point estimates with 95% error bands for the estimates.

Figure 1 depicts the results for four different cases. The top two panels are for the case of a single emission source and the differing levels of preferential sampling. The bottom two show the corresponding results for the multi-city scenario. In each case, the black lines show

the (true) average at each time t over the finite population of 1000 sites. The red lines give the (biased) summaries for each time point together with an indication of the variability over the multiple dataset through the 95% error bars (red dotted lines). Green lines show the HT adjusted summaries at each time point (with dotted lines signifying the associated variation). The adjusted values are extremely close to the true values, a fact reflected in the green lines overlaying the black (which are not visible) in the upper panels.

The overall pattern in all cases is that levels decline over time, reflecting a feature of the superpopulation distribution, and error bands increase in width reflecting smaller samples. When preferential sampling is weak (the two left hand panels) the red curve is relatively close to the black curve although significantly higher. The difference in the right hand panels, when preferential sampling is strong, is much more marked. The HT adjustment improves the estimators of the annual mean in all cases and that improvement is dramatic with strong preferential sampling.

To provide a comparison with the estimates arising from applying the method proposed here, we briefly consider the effects of two much simpler approaches. Given the set of sites that were present at the beginning of the period of interest, when a specific site ceased to be part of the network then the resulting missing data from that time point might be imputed. A regression line might be fit to the available measurements from that site and used to predict measurements from that time. The average for each year would then be calculated using a mixture of observed and predicted measurements for all sites at each year. An even simpler approach might be to fill in missing data for each site using the last recorded measurement, i.e. filling in the missing data with the most recent measurement available for that site. The first approach would require a reasonable amount of data points for each site (here chosen to be five) and is likely to result in negative predictions over the latter time years where the initial decline observed at a site was strong. In such cases, predictions here are truncated at zero. In both case, applying these two simple approaches resulted in the estimates of the yearly averages being overestimated in all four of the cases presented in Figure 1. Filling in the missing values using the last value proved to be consistently higher than the true values with the error increasing over time as might be expected. In the examples shown in the panels of Figure 1, the overestimates in the last (25th) year were 16%, 14%, 8% and 11% for the mild / strong point source and mild /strong multi city cases respectively. The regression approach also consistently overestimated the yearly averages with the corresponding overestimates in the last year being 34%, 18%, 1% and 7%. In the third case (of mild preferential sampling from multiple cities) the regression approach seems to do well as the downward slope of concentrations is close to linear however this is not repeated when there is strong preferential sampling and, as might be expected, applying these simple approaches would result in summaries being overestimated, as would be seen when just using the available data (black lines in Figure 1).

When the site selection probabilities are proportional to size the HT estimator has another desirable feature, that the $\{y_{tu}/\pi_{tu}\}$ ratios will be nearer constant than if the $\{\pi_{tu}\}$ are all equal, as is the case with the unadjusted estimates [Stehman and Overton, 1994]. This will lead to a decrease in the variability associated with the adjusted estimates compared to that for the unadjusted ones, resulting in narrower 95% error bars. In practice, this will be additionally be affected by factors such as the strength of preferential sampling, it's cumulative effect over time and the underlying variability of the population parameters. This can be seen in Figure 1 with the narrow error bars for the adjusted estimates (compared to their unadjusted counterparts) increasing in width over time and with the extent of the variability in the underlying field.

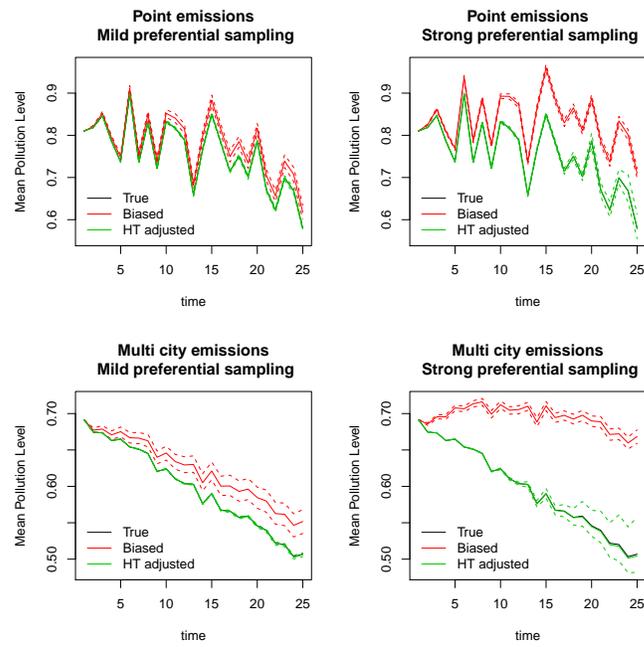


FIG 1. Results from the simulation study of a network that is reduced in size over time. Lines represent the population average over all sites under the super-population models (black), the unadjusted estimates (red) of that mean and the Horvitz-Thompson adjusted estimates (green). Dotted lines show 95% error bands based on 1000 simulated datasets. Upper and lower panels show results for single and multiple point emission scenarios respectively. Note that for most of the times the black line is not visible as it is overlaid by the green line due to the closeness of the adjusted estimates to the true values.

For the second mode of sampling the overall setting is the same as the first, but now the unobserved responses need to be imputed. A GRF superpopulation model is assumed with Model M_1 being $Y_{jtu} = \nu_{jtu} + \epsilon_{jtu}$ where $\nu_{jtu} = \xi_1 x_{jtu} + \xi_2 z_{jtu}$. The structure of Model M_2 is the same with the x replaced by p (as in Equations (5.24) and (5.25)). The preferential sampling covariates $\{z_{jtu}\}$ are specified below. After fitting the $\{\xi_i, i = 1, 2\}$ and the parameters of the covariance models, unobserved y_{jtu} can be imputed by standard geostatistical (or other) methods. Analysis proceeds as follows for each of the two models.

(ii) Expanding adaptive network design

1. time 1: Compute the HT estimate of the population mean $\hat{\mu}_1$ using Equation (4.16). Impute the unobserved responses \hat{y}_{1u} , $u \in \mathcal{D} \setminus S_1$.
2. time 2: Use the imputed responses at time 1 and logistic regression as in Equation (4.18) to estimate the conditional selection probabilities $\hat{\pi}_{2u}$, $u \in S_2$. Then estimate their unconditional probabilities of selection $\hat{\pi}_{2u}^*$ using Equation (4.17). Thus at time 2, for all $u \in S_2$, $\hat{\pi}_{2u}^* = \hat{\pi}_{1u}^* + (1 - \hat{\pi}_{1u}^*)\hat{\pi}_{2u}$. Compute the HT estimate for each of the two scenario models (M_j , $j = 1, 2$):

$$\hat{\mu}_2 = \sum_{u \in S_2} \frac{y_{2u}}{N \hat{\pi}_{2u}^*}$$

3. time 3: Compute z_{2u} using Equation (4.20). Impute the unobserved \hat{y}_{2u} , $u \in \mathcal{D} \setminus S_2$ as described above. Estimate the conditional selection probabilities $\hat{\pi}_{3u}$ for sites in S_3 and then their unconditional selection probabilities $\hat{\pi}_{3u}^*$. Compute the HT estimate for each of the two scenario models (M_j , $j = 1, 2$).

$$\hat{\mu}_2 = \sum_{u \in S_3} \frac{y_{3u}}{N \hat{\pi}_{3u}^*}$$

4. time $t = 4, \dots, T$: Repeat Step 3 after recursively updating it.
5. Repeat the previous steps for each of the 1000 replicate datasets, at each time calculating the HT population average estimates along with their 95% error bands.

Figure 2 shows the corresponding results to those seen in Figure 1 for the expanding network case. As with the decreasing network case, clear differences can be seen between the true finite population means at each time point and those estimated from the data arising from the preferential samples, with the differences being greater in the case of strong preferential sampling. This bias is again markedly reduced when using the HT adjusted estimates. The width of the error bands for both estimates are initially large due to the small sample sizes and they decline in width as time goes on due to the increasing sample sizes. In all cases those for the HT estimator are the narrower of the two for the reasons given above. Perhaps surprisingly this reduction is observed in all four cases over all time, despite the need in this case to impute unobserved responses in the case of the H-T estimator. Imputation does add uncertainty however, as we see in comparison to the widths seen in the the decreasing network example as would be expected given that these bands reflect both variation due to the preferential resampling and that associated with the imputation of unobserved responses.

6. Case study: Black Smoke in the United Kingdom. In this case study we aim to address one of the paper's primary aims, that of adjusting population level estimates of BS levels in the UK over an extended period to make them unbiased. Shaddick and Zidek [to appear]

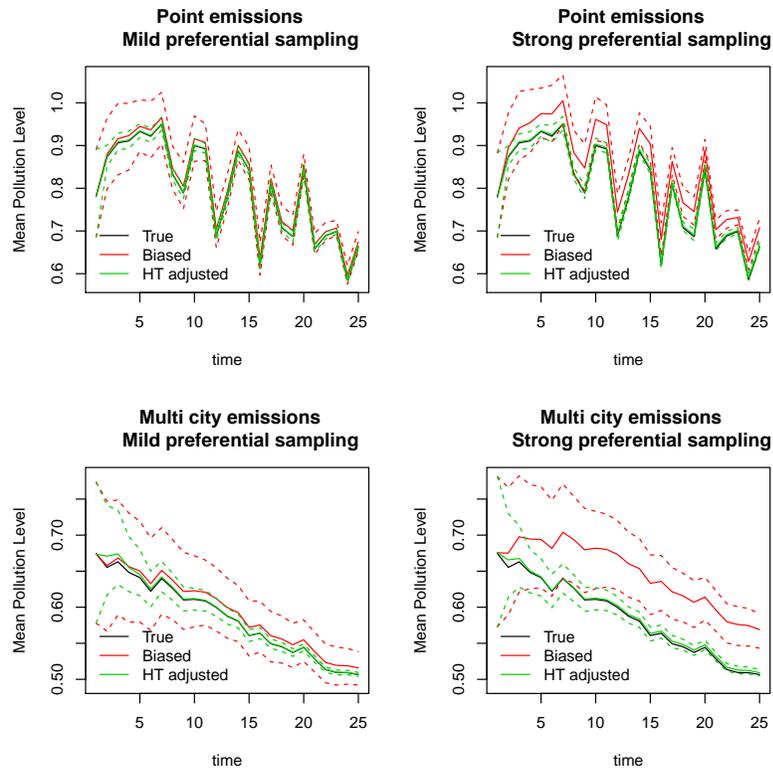


FIG 2. Results from the simulation study of a network that is increasing in size over time. Lines represent the population average over all sites under the super-population models (black), the unadjusted estimates (red) of that mean and the Horvitz-Thompson adjusted estimates (green). Dotted lines show 95% error bands based on 1000 simulated datasets. Upper and lower panels show results for single and multiple point emission scenarios respectively. Note that the black line is often not visible as it is overlaid by the green line due to the closeness of the adjusted estimates to the true values.

provides evidence of preferential sampling in the the reduction in the network over time and here we use the methods developed in Sections 3 and 4 and demonstrated in the simulation studies 5 to adjust estimates of overall average annual concentration levels as well as number of sites out of compliance.

We begin with a summary of the monitoring program for BS in Great Britain. Although air pollution has been a concern for many centuries it became a global health issue in the early parts of the last century after a number of high pollution episodes were linked to increased health risks (Firket [1936]; Ciocco and Thompson [1961]; Ministry of Public Health 1954). As a result, attempts were made to measure air pollution concentrations in a regular and systematic way.

Daily average BS has been shown to be a reasonable predictor of PM_{10} . In general, PM_{10} concentrations are usually higher than BS except during high episodes, and hence, if BS exceeds the PM_{10} limit, it is likely that PM_{10} will also be out of compliance [Muir and Laxen, 1995]. Black smoke (BS) is one of a number of measures of particulate matter, other examples including the coefficient of haze (CoH), total suspended particulates (TSP), as well as direct measurements of PM_{10} , and $PM_{2.5}$. Each of these has been associated with adverse health outcomes (for PM_{10} , Samet et al. [2000]; for $PM_{2.5}$, Goldberg et al. [2001]; for TSP, Lee and Hirose [2010]; for BS, Verhoeff et al. [1996]; for CoH, Gwynn et al. [2000]). Attempts have been made to standardize the measures of pollution by converting the measurements into ‘equivalent’ amounts of PM_{10} , for example $PM_{10} \approx 0.55$ TSP, $PM_{10} \approx \text{CoH}/0.55$, $PM_{10} \approx \text{BS}$ and $PM_{10} \approx PM_{2.5}/0.6$ [Dockery and Pope, 1994].

In 1961 the worlds first co-ordinated national air pollution monitoring network was established in the UK using BS and sulphur dioxide monitoring sites at around 1200 sites. As levels of BS pollution have declined, the network has been progressively rationalized, reduced, moved, replaced and by the mid-nineties it comprised of ca. 200 sites.

The data on annual concentrations of BS (μ g/m³) were used in this case study were obtained from the UK National Air Quality Information Archive. We use data from 1970-1996 and restrict to the case where sites were withdrawn from the network over time. A number of sites were added during this period, but they are almost exclusively ones which reported even higher concentrations, suggesting they were added preferentially. For clarity, we consider only the reduction in the network. We use the 624 sites that were operational in 1970 and which had at least 5 measurements in the following 25 years, and these sites define the finite population, i.e. the concentrations measured at these sites as characterizing the BS field over the UK. For each year, t , data are available from n_t sites, $t = 1, \dots, 26$. Measurements, Z_{it} , are the log of the annual means of the 24 hour mean concentrations of BS divided by a normalization constant to make them unitless (to be able to apply logarithms). Over the study period, the number of sites was reduced from $n_1 = 624$ to $n_{26} = 193$ with the yearly means over all sites, $\sum_{i=1}^{n_i} Z_{it}/n_i$, falling from 60.5 to 9.3 μmg^{-3} over the same period. However the preferential selection used to reduce the network, and demonstrated in Shaddick and Zidek [to appear], suggests the latter number (the sample average of the values the surviving sites) is too high. That calls for an adjustment of the form now described.

Extensive analysis of these data suggest a log–Gaussian random field superpopulation model [Shaddick and Zidek, to appear] for BS, because in addition to the desirable properties of right-skew and non-negativity, there is justification in terms of the physical explanation of atmospheric chemistry [Ott, 1990]. We therefore take the logarithms of BS concentrations after normalizing to eliminate their units of measurement [Monk and Munro, 2010], $Y_{it} = \log(Z_{it}/78)$, where 78 μgm^{-3} represents roughly the average level of BS at the beginning of the time at which the

network was operational.

Working on the log-scale described above, we applied the methods described in Section 5 for (i) decreasing networks to adjust the annual arithmetic averages for the effects of preferential sampling. Two characteristics associated with the responses seem of natural interest. The first we consider is the set of annual averages across these 624 sites as these could be published to show the effect of regulatory policy over time. The left hand panel of Figure 3 shows the estimated geometric annual mean levels over time (dotted black line) together with the HT adjusted ones (solid grey line). It clearly shows the adjustment reduces the estimates of the average levels. Since the standard unit for calculating relative risks of particulates in health effects analysis is $10\mu\text{gm}^{-3}$ the difference seems important, being more than one of these standard units over much of the period. This is because, on the log-scale, responses with low values get high weights in the HT adjusted arithmetic average, since their chances of making the cut in every successive year, say to 1985 for example is very small. As in Section 5, we also briefly consider the effects of two very simple approaches to ‘filling in’ the missing data after a site has been excluded from the network; (i) using the last recorded value throughout the following years and (ii) using prediction from linear regression. Further details can be found in Section 5. As might be expected, using the first approach here results in much higher estimates of the annual averages; for example in 1996 the estimate was $19.1\mu\text{gm}^{-3}$ compared to $9.8\mu\text{gm}^{-3}$ obtained from the available data and $1.1\mu\text{gm}^{-3}$ using the HT approach. The corresponding estimate using the linear regression approach was $2.4\mu\text{gm}^{-3}$. Whilst this appears a not unreasonable estimate, the simplistic nature of the correction means that in the the early years of the analysis it produces estimates which are much higher than that obtained from the available data and the HT estimates, in 1972 for example the estimated annual average would be $60.5\mu\text{gm}^{-3}$ using linear regression compared to $50.6\mu\text{gm}^{-3}$ using the available data and $48\mu\text{gm}^{-3}$ the HT approach respectively.

The second characteristic we consider is potentially of even greater operational importance; the number of sites in non-attainment, i.e. those which do not comply with of air quality standards in given year. This number is a surrogate for the cost of mitigation for putting the BS concentrations into compliance. For example, as part of the analysis of the impact of the various ozone standards considered by the EPA’s CASAC Ozone Committee in 2008, the EPA Staff predicted the fraction of monitored counties in the United States that would be out of compliance. For the standards that were finally proposed by the Committee, that percent was found to be 86%. Although the US Clean Air Act (epa.gov/oar/caa/title1.html) of 1970, under whose mandate the CASAC was created, rules out economic impact in consideration of standards designed to protect public health, nevertheless policy making cannot ignore the cost of attainment which can be substantial. The right hand panel of Figure 3 shows the number of sites each year that exceed the 1980 EU guide value of $34\mu\text{gm}^{-3}$ [European Commission, 1980]. The dotted black line is the number of exceeding sites based on the recorded data with the solid grey line the numbers after adjustment for the preferential sampling. The unadjusted numbers are the fraction in the monitoring network out-of-compliance multiplied by the finite population total of $N = 624$. Their adjusted counterparts are found by applying the HT weights to the 1s present in the summation used to calculate that fraction. Table 3 shows the number of sites exceeding the EU limit of $68\mu\text{gm}^{-3}$ and the guide values of 51 and $34\mu\text{gm}^{-3}$ where a reduction of the number of exceedances can be seen when adjusting for the preferential sampling. For example, in 1974, the crude estimate gives 211 of the 624 sites out of compliance with the $34\mu\text{gm}^{-3}$ criterion, while its adjusted counterpart is just 189. This is a substantial difference, given the economic cost of mitigation. Note the large number of zeros reflect the decline over time in the levels of BS.

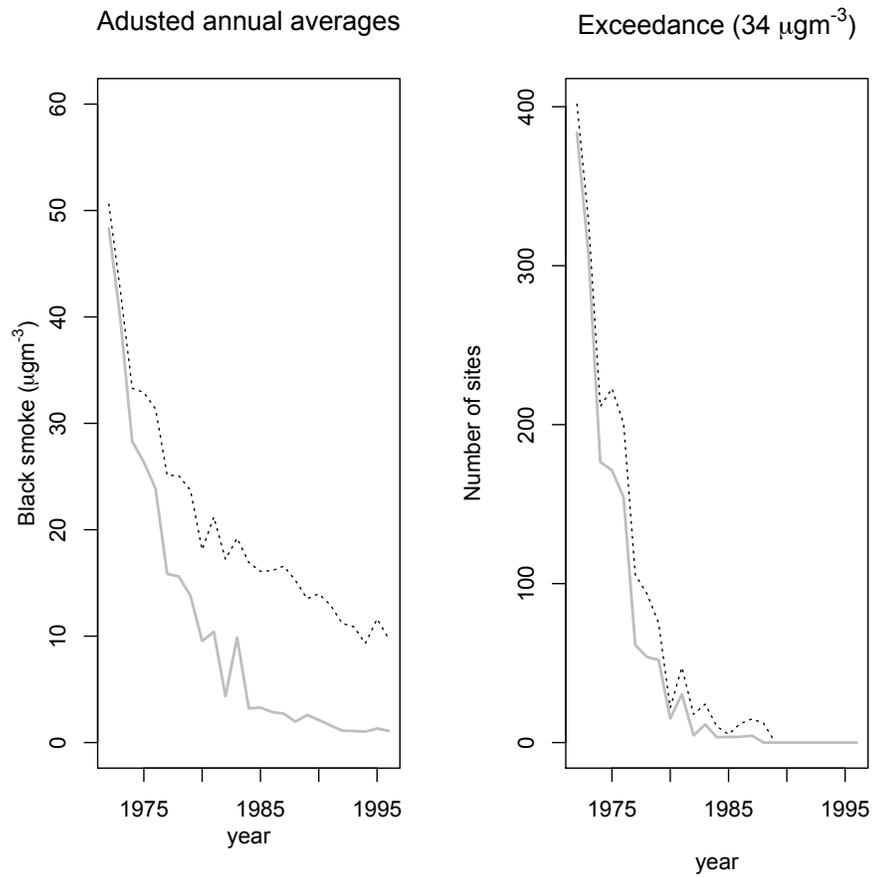


FIG 3. Changes in the levels of black smoke within the UK from 1970-1996 and the effects of adjusting estimates of annual indices. The left hand panel shows the annual means over all sites (dotted black line) together with adjusted values (solid grey line). The right hand panel shows the number of sites exceeding the EU guide value of $34 \mu\text{gm}^{-3}$ (dotted black line) together with the corrected values (solid grey line).

Overall we see quite a substantial overestimation of important finite population parameters due to the preferential sampling in published estimates.

TABLE 3

Estimated number of sites exceeding regulatory guide values for black smoke, with and without adjusted for preferential sampling using Horivitz–Thompson estimators.

	Limit 68 μgm^{-3}		Guide 51 μgm^{-3}		Guide 34 μgm^{-3}	
	Unadjusted	Adjusted	Unadjusted	Adjusted	Unadjusted	Adjusted
1972	129	123	236	225	402	384
1973	73	68	153	143	327	306
1974	31	28	94	84	211	189
1975	21	19	58	52	223	201
1976	19	18	50	47	201	189
1977	7	6	23	20	106	94
1978	7	7	18	17	94	87
1979	8	7	21	19	75	71
1980	0	0	6	6	22	21
1981	2	2	11	10	47	42
1982	0	0	0	0	18	9
1983	0	0	10	5	24	16
1984	0	0	0	0	10	8
1985	0	0	0	0	5	4
1986	0	0	0	0	12	6
1987	0	0	0	0	15	7
1988	0	0	0	0	12	3
1989	0	0	0	0	0	0
⋮	⋮	⋮	⋮	⋮	⋮	⋮
1996	0	0	0	0	0	0

7. Discussion. A number of methods have been discussed in this paper for modeling the probability of selection in preferential sampling and we have developed a general framework using a superpopulation modeling approach. Taking a public policy perspective, we have emphasized the HT–approach to mitigating the effects of preferential sampling in order to get unbiased estimators.

Having the space–time series of sites in Section 6 enables us to do something that is not possible with only spatial data, namely to study the preferential selection process itself. This is performed using logistic regression to estimate the selection probabilities (Sections 5 and 6). The results from the simulation studies in Section 5 suggest the method proposed in this paper compensates for the effect of preferentially sampling and reduces bias. Section 6 shows that this adjustment can be substantial in practice.

The case study in Section 6 demonstrated the use of the method where the parameters being estimated are numerical features of the finite population of exposures of air pollution. The finite population in this case were monitoring sites in the UK that were measuring BS in 1970. From that time there was a dramatic reduction in the size of the monitoring network and subsequently only subsets of those exposures were measured. In the case study we apply the methods we have developed to adjust these estimates for the effects of preferential sampling. The results show reductions in the estimates, illustrating how the preferential siting of monitors where exposures are high gives an exaggerated impression of the level of BS and the number of sites that are in noncompliance. Note that selection bias can accumulate over time and can require increasingly greater adjustments. Although the effect on estimates of the number of sites out of compliance

appears not to be as dramatic, it is substantial, especially considering that forcing attainment of standards can entail large costs.

In the case study, we assumed a log-Gaussian superpopulation model for BS and so the geometric, rather than arithmetic annual average, must be used to characterize the finite population's annual mean level. This metric is often used for particulate air pollution, see for example Mueller [1994]. In Section 3 (Case 2) we show that in such cases spatial-correlation-adjusted (unequal) weights ought to be used to characterize the finite population mean levels. However this is extremely unlikely to be the case in reality and as a concession to standard practice we use equal weights, leaving this issue for exploration in future work. Asymptotic theory and variance approximations are available for the Horvitz-Thompson estimator, which would enable approximate error bands to be computed for the adjusted estimates. However these would not include the additional uncertainty associated with the estimation of the selection probabilities themselves. So we plan in future work to develop an approach that would also involve the construction of measures of uncertainty associated with the sampling weights and consequently the adjusted annual averages and exceedances for a single data analysis (in the simulation study 95% error bars are obtained by repeated simulation). If both the modelling of the weights and the adjustment process were combined within a Bayesian framework it may be possible to propagate the uncertainty in the estimation of the weights (including that which may arise due to spatial prediction in the case of expanding networks) through to the adjusted values.

The HT-approach will not always be the most appropriate approach. In some cases a likelihood based approach may be feasible provided that the preferential sampling can be modelled. That could be the case in the context of air pollution and health in epidemiological analyses, for as Guttorp and Sampson [2010] points out, the air pollution monitoring sites may be intentionally located for reasons such as the need to measure: (i) background levels outside of urban areas; (ii) levels in residential areas; and (iii) levels near pollutant sources. Then the method in D10 may be useful. A possible alternative approach to the method in D10 is described in Zidek and Shaddick [2012]. While it resembles the point process model, it is designed for discrete site domains as seen in this paper. Moreover, being based on intensities for paralyzable particle counters, it would allow for preferential sampling designs with varying intensities over time.

The approach taken in the paper will work best when at least some representatives of the general population of sites continue to be sampled since the selection weights would then be able to compensate for their under representation in computing population statistics. In the case study, the BS network was originally set up to try and provide monitoring for a cross-section of expected cities and pollution levels although in reality the original set of locations would have included some form of preferential sampling. In the absence of such representation or good background knowledge of how the biased selection was made, there would seem to be no alternative but to augment the network with some, possibly temporary monitors. That leads to a design problem about the optimal selection of those sites. In this case, then we would be unbiassing the design rather than unbiassing the estimates, and that would need a different approach than that described in this paper.

The analyses reported in this paper and its predecessors (Diggle et al. [2010], Gelfand et al. [2012], Pati et al. [2011]) coupled with the importance and wide spread use of environmental monitoring networks, points to the need for further exploration and confirmation of the results of these analyses. We recognize there are limitations in the preferential sampling models used in the simulation studies in these papers. In practice site selection is complex involving committees, guidelines and negotiations, and local administrators in affluent areas demanding a monitor in

their municipality. However, they do convincingly demonstrate that preferential site selection does have serious adverse consequences, something that does not appear to have been recognized by agencies charged with formulating regulatory guidelines. Network data are commonly used as if they represent a true reflection of underlying environmental fields. Of course, there are occasions when sites must be preferentially located, e.g. to check adherence to standards around industrial sources. In such cases it may be undesirable for the data to be used for other, possibly unintended, purposes but if it is so used it should be adjusted for the possible effects of preferential sampling using methods such as that presented in this paper.

Acknowledgements. The genesis of the work described in this paper was the 2009-10 thematic program on spatial analysis for environmental mapping, epidemiology and climate change developed by the Statistical and Applied Mathematics Institute on spatial statistics. We are grateful to members of its Working Group on preferential sampling, of which the second author was a member, for discussions that stimulated our interest in that topic. We would like to thank Yi Liu for his assistance in the comparison of methods. We would also like to offer great thanks to the Editor, Associate Editor and anonymous referees who offered valuable suggestions which have greatly improved the paper, one of which offered an extraordinarily detailed list of comments and corrections without which the revision of the manuscript would have been a much greater task.

References.

- B. Ainslie, C. Reuten, DG Steyn, N.D. Le, and J.V. Zidek. Application of an entropy-based Bayesian optimization technique to the redesign of an existing monitoring network for single air pollutants. *Journal of environmental management*, 90(8):2715–2729, 2009.
- D. A. Binder. On the variances of asymptotically normal estimators from complex surveys. *International Statistical Review/Revue Internationale de Statistique*, pages 279–292, 1983.
- D.A. Binder and P. Zdenek. Use of estimating functions for estimation from complex surveys. *J Amer Statist Assoc*, 89:1035, 1043 1994.
- H. Chang, A.Q. Fu, N.D. Le, and J.V. Zidek. Designing environmental monitoring networks to measure extremes. *Environmental and Ecological Statistics*, 14(3):301–321, 2007.
- Giuseppe Cicchitelli and Giorgio E Montanari. Model-assisted estimation of a spatial population mean. *International Statistical Review*, 80(1):111–126, 2012.
- A. Ciocco and D.J. Thompson. A follow-up of donora ten years after: methodology and findings. *Am J Public Health Nations Health*, 51:155–164, 1961.
- P. Dawid. Discussion of Geostatistical inference under preferential sampling by Diggle, P.J., Menezes, R. and Su, T. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 59(2):191–232, 2010.
- P.J. Diggle, R. Menezes, and T. Su. Geostatistical inference under preferential sampling. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 59(2):191–232, 2010.
- D. Dockery and CA III Pope. Acute respiratory effects of particulate air pollution. *Annu. Rev. Public Health*, 15: 107–132, 1994.
- EPA. Air quality criteria for ozone and related photochemical oxidants. 2006. EPA/600/R-05/004aF-cF.
- EPA. National Lakes Assessment: a collaborative survey of the nation’s lakes. 2009. EPA 841-R-09-001.
- European Commission. Council directive 80/779/EEC of 15 July 1980 on air quality limit values and guide values for sulphur dioxide and suspended particulates. 1980.
- J. Firket. Fog along the Meuse valley. *Trans Faraday Soc*, 32:1191–1194, 1936.
- Alan E Gelfand, Sujit K Sahu, and David M Holland. On the effect of preferential sampling in spatial prediction. *Environmetrics*, 23(7):565–578, 2012.
- VP Godambe and M.E. Thompson. Parameters of superpopulation and survey population: their relationships and estimation. *International Statistical Review/Revue Internationale de Statistique*, pages 127–138, 1986.
- M.S. Goldberg, R.T. Burnett, J. C. Bailar, 3rd, R. Tamblin, P. Ernst, J. Flegel, K. Brook, Y. Bonvalot, R. Singh, M. F. Valois, and R. Vincent. Identification of persons with cardiorespiratory conditions who are at risk of dying from the acute effects of ambient air particles. *Environ Health Perspect*, 109(Suppl 4):487–494, 2001.
- P. Guttorp and P. Sampson. Discussion of Geostatistical inference under preferential sampling by Diggle, P.J., Menezes, R. and Su, T. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 59(2):191–232, 2010.

- R.C. Gwynn, R. T. Burnett, and G.D. Thurston. A time-series analysis of acidic particulate matter and daily mortality and morbidity in the Buffalo, New York, region. *Environ Health Perspect*, 108:125–133, 2000.
- D.G. Horvitz and D.J Thompson. A generalization of sampling without replacement from a finite universe. *J Amer Statist Assoc*, 47:663–685, 1952.
- Itai Kloog, Francesco Nordio, Brent A Coull, and Joel Schwartz. Incorporating local land use regression and satellite aerosol optical depth in a hybrid model of spatiotemporal PM_{2.5} exposures in the Mid-Atlantic states. *Environmental Science & Technology*, 46(21):11913–11921, 2012.
- JF Lawless, JD Kalbfleisch, and CJ Wild. Semiparametric methods for response-selective and missing data problems in regression. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 61(2):413–438, 1999.
- N.D. Le and J.V. Zidek. *Statistical analysis of environmental space-time processes*. Springer Verlag, 2006.
- A. Lee and Y. Hirose. Semi-parametric efficiency bounds for regression models under response-selective sampling: the profile likelihood approach. *Annals of the Institute of Statistical Mathematics*, 62(6):1023–1052, 2010.
- Kung-Yee Liang and Scott L Zeger. Longitudinal data analysis using generalized linear models. *Biometrika*, 73(1): 13–22, 1986.
- P. Monk and L.J. Munro. *Maths for Chemistry: A chemists toolkit of calculations*, Second Edition. Oxford University Press, 2010.
- Stephen F Mueller. Characterization of ambient ozone levels in the Great Smoky Mountains National Park. *Journal of Applied Meteorology*, 33(4):465–472, 1994.
- D. Muir and D.P.H. Laxen. Black smoke as a surrogate for PM₁₀ in health studies? *Atmospheric Environment*, 29 (8):959–962, 1995.
- W Ott. A physical explanation of the lognormality of pollutant concentrations. *Journal of the Air & Waste Management Association*, 40(10):1378–1383, 1990.
- Debdeep Pati, Brian J Reich, and David B Dunson. Bayesian geostatistical modelling with informative sampling locations. *Biometrika*, 98(1):35–48, 2011.
- D. Pfeffermann. The role of sampling weights when modeling survey data. *International Statistical Review/Revue Internationale de Statistique*, pages 317–337, 1993.
- JNK Rao, AJ Scott, and C.J. Skinner. Quasi-score tests with survey data. *Statistica Sinica*, 8:1059–1070, 1998.
- J.M. Samet, F. Dominici, F.C. Curriero, I. Coursac, and S. L. Zeger. Fine particulate air pollution and mortality in 20 U.S. cities, 1987–1994. *N Engl J Med*, 343:1742–1749, 2000.
- C.-E. Sarndäl, B. Swensson, and J. Wretman. *Model assisted survey sampling*. Statistics. Springer New York, 2003.
- Peter Schumacher and James V Zidek. Using prior information in designing intervention detection experiments. *The Annals of Statistics*, pages 447–463, 1993.
- A.J. Scott and C.J. Wild. Fitting binary regression models with response-biased samples. *Canadian Journal of Statistics*, 39(3):519–536, 2011.
- G. Shaddick and J.V. Zidek. A case study in preferential sampling: long term monitoring of air pollution. *Spatial Statistics*, to appear.
- Stephen V Stehman and W Scott Overton. Comparison of variance estimators of the Horvitz-Thompson estimator for randomized variable probability systematic sampling. *Journal of the American Statistical Association*, 89 (425):30–43, 1994.
- A. P. Verhoeff, G. Hoek, and J.H. Schwartz, J.; van Wijnen. Air pollution and daily mortality in Amsterdam. *Epidemiology*, 7:225–230, 1996.
- James V Zidek, Weimin Sun, and Nhu D Le. Designing and integrating composite networks for monitoring multivariate Gaussian pollution fields. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 49(1): 63–79, 2000.
- J.V. Zidek and G. Shaddick. Unbiasing estimates from preferentially sampled spatial data. Technical Report 268, University of British Columbia, 2012.

DEPARTMENT OF STATISTICS
 UNIVERSITY OF BRITISH COLUMBIA
 2207 MAIN MALL VANCOUVER, BC
 V6T 1Z4, CANADA
 E-MAIL: jim@stat.ubc.ca

DEPARTMENT OF MATHEMATICAL SCIENCES
 UNIVERSITY OF BATH
 BATH, BA2 7AY
 UK
 E-MAIL: masgs@bath.ac.uk

DEPARTMENT OF STATISTICS
 UNIVERSITY OF BRITISH COLUMBIA
 2207 MAIN MALL VANCOUVER, BC
 V6T 1Z4, CANADA
 E-MAIL: cgtaylor@stat.ubc.ca