U.S. DEPARTMENT OF ENERGY

PNNL-23244

Prepared for the U.S. Department of Energy
under Contract DE-AC05-76RL01830

Application of the Combined Judgmental and Randomly Placed Sampling Method

SL Morris
BG Amidan
LH Sego

March 2014

Pacific Northwest
NATIONAL LABORATORY

Proudly Operated by Battelle Since 1965

PACIFIC NORTHWEST NATIONAL LABORATORY
operated by
BATTELLE
for the
UNITED STATES DEPARTMENT OF ENERGY
under Contract DE-AC05-76RL01830

Printed in the United States of America

Available to DOE and DOE contractors from the
Office of Scientific and Technical Information,
P.O. Box 62, Oak Ridge, TN 37831-0062;
ph: (865) 576-8401
fax: (865) 576-5728
email: reports@adonis.osti.gov

Available to the public from the National Technical Information Service
5301 Shawnee Rd., Alexandria, VA 22312
ph: (800) 553-NTIS (6847)
email: orders@ntis.gov <http://www.ntis.gov/about/form.aspx>
Online ordering: http://www.ntis.gov

This document was printed on recycled paper.
(8/2010)

PNNL-23244

Application of the Combined Judgmental and Randomly Placed Sampling Method

SL Morris
BG Amidan
LH Sego

March 2014

Prepared for the U.S. Department of Homeland Security
Science and Technology Directorate
and the U.S. Department of Energy
under Contract DE-AC05-76RL01830

Pacific Northwest National Laboratory
Richland, Washington 99352

Contents

1.0 Introduction

In the wake of the 2001 anthrax incidents in the United States, there has been considerable discussion about the role of judgmental and probabilistic sampling strategies, particularly in the context of clearance sampling. The Combined Judgmental and Random (CJR) (Sego et al. 2010) methodology provides a statistical methodology that allows practitioners to leverage information from both judgmental and randomly selected sample locations and then quantify the confidence in negative results. The objective of the CJR sampling design is to demonstrate, with high probability, that a high percentage of the decision area (or population) is acceptable, given that *none of the samples are unacceptable*. The purpose of this document is to assist practitioners to appropriately use the CJR methodology and to implement this sampling design in Visual Sample Plan (VSP) (VSP Development Team, 2014).

We begin by defining some fundamental terms. A *decision area* is the geographical or spatial area (indoor or outdoor) about which practitioners wish to draw a conclusion on the basis of sampling and other relevant information. Decision areas typically include all or part of a building, but they may be any contaminated (or potentially contaminated) area for which characterization or clearance is required. In this report, a *statistical sampling* approach refers to any sampling approach based on a mathematical or statistical model that provides confidence or probability statements about the decision area. Statistical sampling approaches typically involve taking samples at random locations, or at the nodes of a randomly located grid or lattice that is superimposed on the decision area. The term *probabilistic sampling* is often used to describe statistical sampling designs that exclusively use randomly located samples. This terminology implies that probabilistic statements can only be made based upon the results of randomly placed samples. However, the CJR method utilizes a Bayesian model that gives rise to a formal probabilistic statement regarding the acceptability of the decision area based upon the information obtained from both judgmental and random samples. Judgmental samples play a vital role in the characterization and clearance process because they include the informed prior belief and expert knowledge of personnel involved in the sampling process.

1.1 Example

The following example illustrates a situation in which the CJR approach can be used to create a sampling design that is more efficient than traditional statistical sampling (that uses only randomly located samples). This example will be referred to throughout the text.

A terrorist releases anthrax within a very large building. Following emergency response and law enforcement investigation, teams go in and perform characterization sampling, which is followed by decontamination. A practitioner wants to devise a sampling scheme (using culture methods) to show that a large area within the building is sufficiently safe to reoccupy. A sample result will be designated as “acceptable” if no growth occurs in the culture, and “unacceptable” if not. The practitioner plans to take samples from high-risk locations where contamination might still be present (if any is present at all). However, not all stakeholders are comfortable with only taking those judgmental samples and want to augment the judgmental sampling with some samples obtained from a systematic grid across the entire area. Before the area can be released for re-occupancy, the practitioner must demonstrate they are at least 95% confident that at least 99% of the floor and wall surface achieves the standard of no culturable growth.

The sampling process involves wiping a 1 square foot area and analyzing each wipe for anthrax colony forming units. If anthrax is detected in any sample, then additional sampling will be required to delineate the contaminated area or the entire area will be decontaminated, depending on relative costs.

Because spore strips set out during decontamination showed the decontamination process was effective and previous studies have demonstrated that this process has been successful under similar circumstances, the practitioner has a strong belief the area is now acceptably free of contamination. If they had to quantify that prior belief, they would expect (at most) that no more than 1% of the high-risk area (which will be sampled judgmentally) would have culturable growth. The practitioner has chosen to take 30 judgmental samples in the locations they believe are at least 2 times more likely to still be contaminated than other areas within the building. These judgmental samples are taken from table tops, doorways, and air vents near the release source.

The floor and wall surfaces of the area within the building constitute some 20,000 square feet. How many additional randomly located samples would be required to demonstrate with 95% confidence that 99% of the surface area is acceptable, provided that none of the samples detect contamination?

Figure 1.1.1 shows a screen shot from VSP with the inputs required for implementing CJR sampling for the example. The other tabs in the Presence / Absence Sampling dialog box will not be discussed here, but are discussed in the VSP help documentation. In the example, 39 random samples (along with the 30 judgmental samples) will need to be taken and found to produce no culturable growth in order to claim with 95% confidence that at least 99% of the decision area would result in no culturable growth, if it were sampled. This particular application of CJR sampling is discussed further in Section 2.1.1.

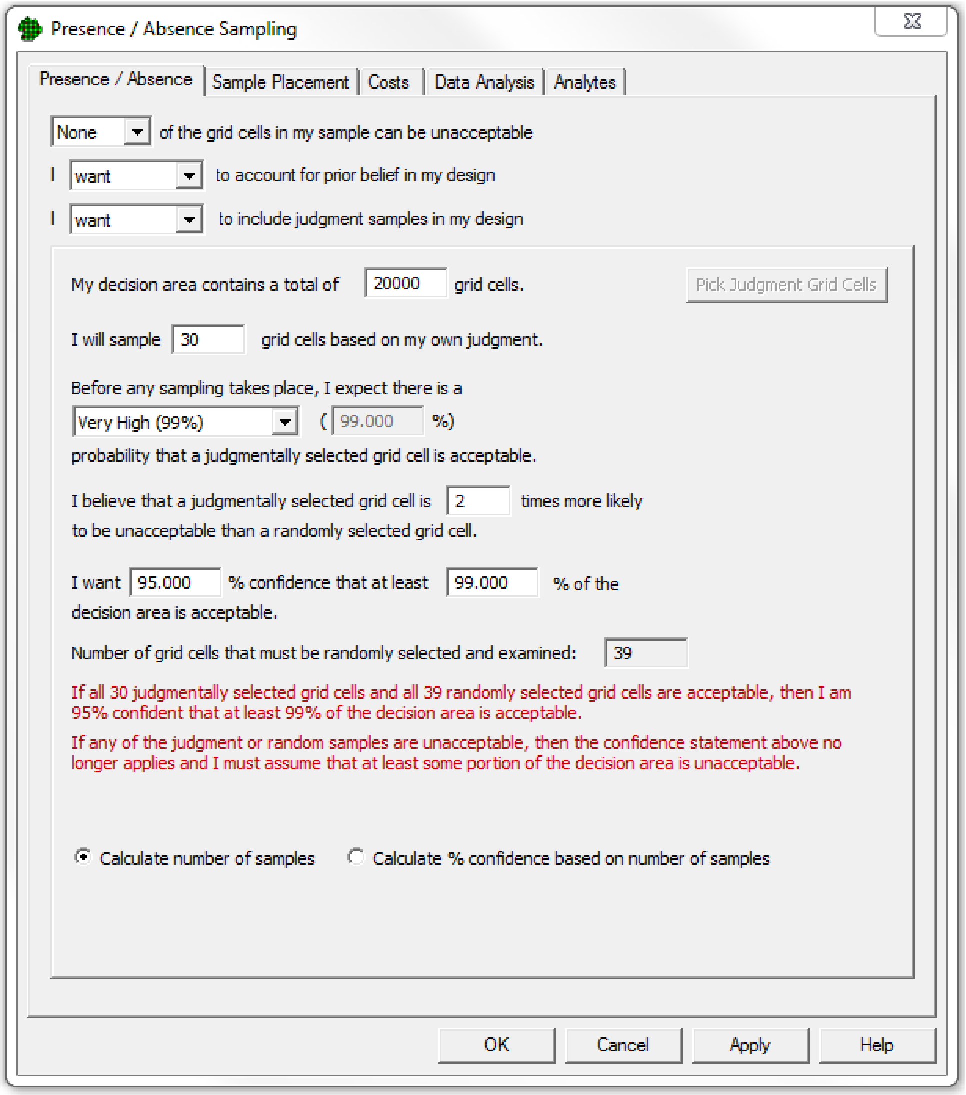

Figure 1.1.1. VSP screen displaying inputs for one type of CJR sampling for the example.

2.0 Methodological Overview

We now briefly describe the CJR methodology in general terms. While the following discussion is presented in terms of sampling conducted within a decision area (such as room or collection of rooms within a building), the CJR methodology is equally applicable to the sampling of any finite population of items—in which case the decision area is analogous to the population of items and the grid cell sampling locations are analogous to the individual items that will be sampled.

Combined Judgmental and Random (CJR) sampling (Sego et al, 2007, 2010) is a statistical method that can be used to determine the number of randomly located samples that should be taken in addition to a predetermined number of judgmental samples to establish with $C \times 100\%$ confidence that at least $\lambda \times 100\%$ of the decision area is acceptable—provided that none of the samples are found to be unacceptable.

The statistical model of the CJR methodology is based on binary outcomes, which we will label as "acceptable" or "unacceptable." Practitioners are required to formally define the criteria that determines whether a sample is acceptable. For example, this criteria may be based on 1) the presence or absence of a particular quality, 2) a quantitative sample result being acceptable or unacceptable as defined by an action level threshold, 3) contamination being detected or not detected, etc.

Sometimes judgmental samples are referred to as "targeted" samples. We use the two terms interchangeably. We define a judgmental sample to be any sample whose location is determined by informed prior belief and expert knowledge and not in a random fashion.

The CJR method requires that all surfaces in the decision area be divided into non-overlapping, equal-size grid cells of specified size that correspond to the sampling methodology, e.g., 1 square foot. While the CJR method was initially designed for use inside buildings, it may be used outdoors if the decision area can be divided into grid cells. In the case of item sampling, a grid cell would be analogous to a specific item.

The size of the grid cell should correspond to the footprint of the sampling methodology (i.e. the area sampled by the swab, wipe or vacuum). If more than one sampling methodology is to be employed in a decision area, the size of the grid cell should be chosen to match the sampling methodology with the smallest footprint. The location of samples that will be taken using methodologies with larger footprints should be assigned in a consistent fashion, e.g. the sample is centered on the smaller grid cell that was assigned by VSP, or the upper-left corner of the larger sample is aligned with the upper-left corner of the grid cell assigned by VSP, etc. Samples with larger footprints would still count as a single sample. While this approach to multiple sampling methodologies is conservative, it ensures that the desired confidence level is preserved.

The CJR design is especially suited for use in decision areas where unacceptable outcomes are deemed unlikely a priori. If at any time during the sampling process, one of the samples is found to be unacceptable, the decision area is declared to be unacceptable and no further samples for the CJR design need be taken. If this occurs, it may be desirable to implement a hot spot or geospatial sampling plan to characterize the extent of the unacceptable locations or items, or the entire space may be decontaminated (again).

We presume that judgmental samples are taken in areas that are more likely to be unacceptable than areas available for random sampling. Consequently, if none of the judgmental sample results are unacceptable, that information can be used to reduce the number of random samples required to achieve the desired confidence level.

The general process of a presence / absence (clearance) sampling plan includes taking samples from the decision area, analyzing the samples, and then, if all the samples are acceptable, making a confidence statement similar to "there is $C \times 100\%$ confidence that at least $\lambda \times 100\%$ of the decision area is acceptable," provided all samples are acceptable. Section 2.1 discusses the methodologies that can be used to determine the clearance sampling plan; Section 2.2 further discusses the inputs needed in VSP for these methods; and Section 2.3 lists the assumptions that are made when applying these methods.

2.1 Presence / Absence (Clearance) Sampling Design

When creating a clearance sampling plan, the practitioner needs to determine their level of confidence, *C*, and the proportion of the decision area, λ, that one desires to show is acceptable. For example, if the practitioner decided on $C = 0.95$ and $\lambda = 0.99$, then they would be 95% confident that at least 99% of the decision area is acceptable. As the desired confidence and/or the desired proportion of acceptable area are increased, the number of samples required to make the statement also increases. In general, the sample size is much more sensitive to λ than *C*.

The number of samples required to make a confidence statement during clearance sampling also depends on incident details and how those samples will be taken. The CJR method is based on the conceptual model that the decision area can be divided into two regions, each with differing risk of contamination. The region with lower risk (called the low-risk area) is largely expected to be acceptable, a priori, and likely consists of a large majority of the total area. The second area (called the high-risk area) is designated by practitioners as an area that is more likely to be contaminated, if any contamination still exists. The location, boundaries, and surface types that comprise the high-risk area are determined by expert judgment. If the high-risk area is fully sampled (these are called judgmental or high-risk samples), then the low-risk area can be randomly sampled (low-risk samples) less intensively than it would have been in the absence of any judgmental sampling.

Two additional parameters concerning the high-risk samples also influence the total number of required samples. These parameters are:

1. P_J – The expected fraction (prior to sampling) of the high-risk area that will be acceptable, and

2. *r* – The likelihood ratio between the high-risk samples being unacceptable compared to the low-risk samples. A high-risk judgmental sample location is, on average, *r* times more likely to be unacceptable than a low-risk random sample location. Effectively, an acceptable judgmental sample will contribute to our confidence *r* times as much as an acceptable random sample.

These parameters are further defined, along with suggested values for P_J and r, in Sections 2.3.2 and 2.3.3. Sections 2.1.1 through 2.2.1 discuss the requirements for applying the method with and without judgmental samples, as well as with or without prior belief. The decision area must be well-defined and the total number of grid cells (the population size) must be determined. In each of these cases, the number of random samples that achieves the desired confidence level can be determined by solving the confidence function using standard numerical root-finding algorithms. Additional details can be found in Sego et al. (2010) and the VSP help files (VSP Development Team, 2014). Recommendations for parameter values can be found in Section 2.2.

2.1.1 Accounting for prior belief & including judgmental samples

To receive the full benefit of the CJR method, a practitioner needs to define the high and low-risk areas, estimate the a priori expected fraction of the high-risk area that will be acceptable, and estimate the likelihood ratio of a high-risk sample being unacceptable compared to a low-risk sample. An example of this application (specifically for the example) is discussed in Section 1.1, and the corresponding VSP settings are given in Figure 1.1.1.

2.1.2 Including judgmental samples but <u>not</u> accounting for prior belief

If prior belief concerning the likelihood of high-risk samples being acceptable is not available, judgmental sampling can still be performed. This is a special case of the CJR method, where P_J is set to 0.50, which corresponds to a uniform, or non-informative prior distribution. Beginning with the same parameter settings in the example, if a practitioner is unable or unwilling to specify the probability that no more than 1% of the high-risk area is still contaminated (hence, we do not include prior belief in our design), then 232 random samples are required, as shown in **Figure 2.1.1**.

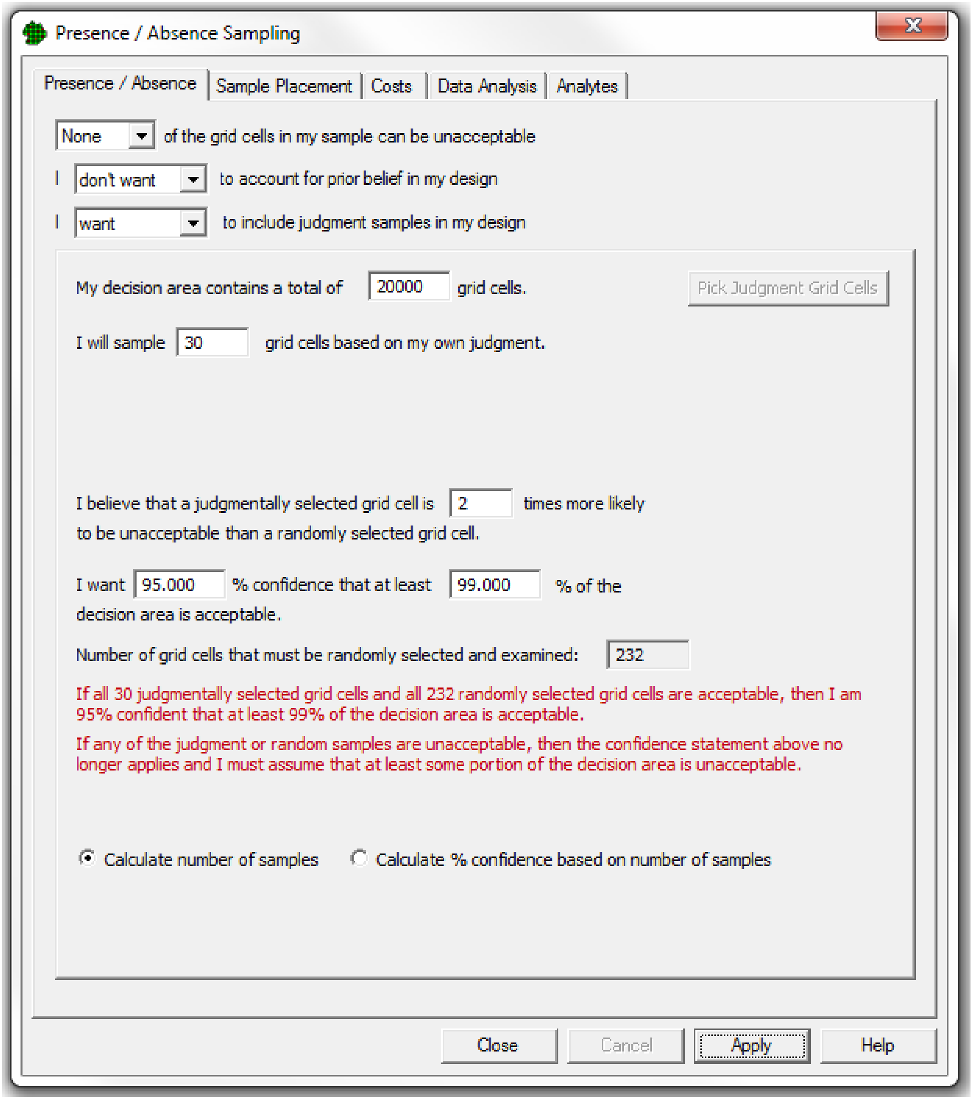

Figure 2.1.1. VSP screen showing the inclusion of judgmental samples but not accounting for prior belief.

2.2 Accounting for prior belief but not including judgmental samples

If a decision area is not divided into high and low-risk areas, then all samples will contribute the same amount of information to the confidence statement. If practitioners have a prior belief concerning the expected fraction of the decision area that will be acceptable, the CJR method can still be applied. This is a special case where the number of judgmental samples is set to 0, *r* is set to 1, and P_J is now the expected fraction of the *entire* decision area that is believed to be acceptable. In the example, if a practitioner wants to account for prior belief in the sampling design but does not want to specify judgment samples, this design is implemented in VSP as shown in **Figure 2.2.1**. In this case, 197 random samples need to be taken.

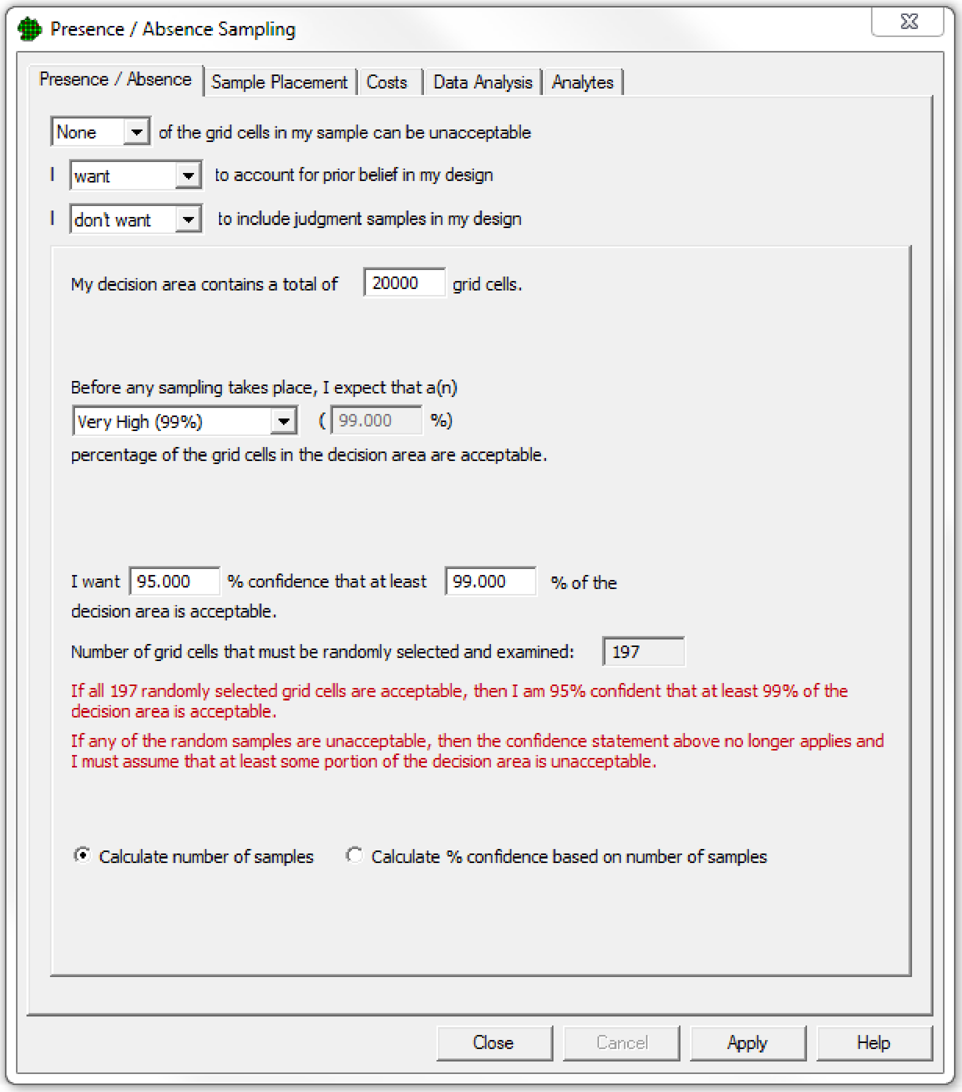

Figure 2.2.1. VSP screen showing prior belief incorporated but not judgmental samples.

2.2.1 Not accounting for prior belief & not including judgmental samples

When practitioners do not wish to account for prior belief and do not want to use judgmental samples, accept-on-zero attribute compliance sampling (AOZ-ACS) (Bowen and Bennett 1988; Squeglia 1994) is appropriate for clearance sampling. The AOZ-ACS method makes the same types of confidence statements as the CJR method [1]. After determining the total number of grid cells and specifying the confidence parameters C and λ, VSP calculates the number of cells that must be sampled and measured. As with the CJR method, if the measurement of one or more of the grid cells is unacceptable, the (C, λ) clearance statement cannot be made. Additional details can be found in the VSP help files (VSP Development Team, 2014). VSP shows that 296 random samples need to be taken in this situation, as shown in **Figure 2.2.2**.

[1] Strictly speaking, the CJR method gives a Bayesian probability statement, and the AOZ-ACS method gives a frequentist confidence statement. In practice, both confidence statements are interpreted the same way.

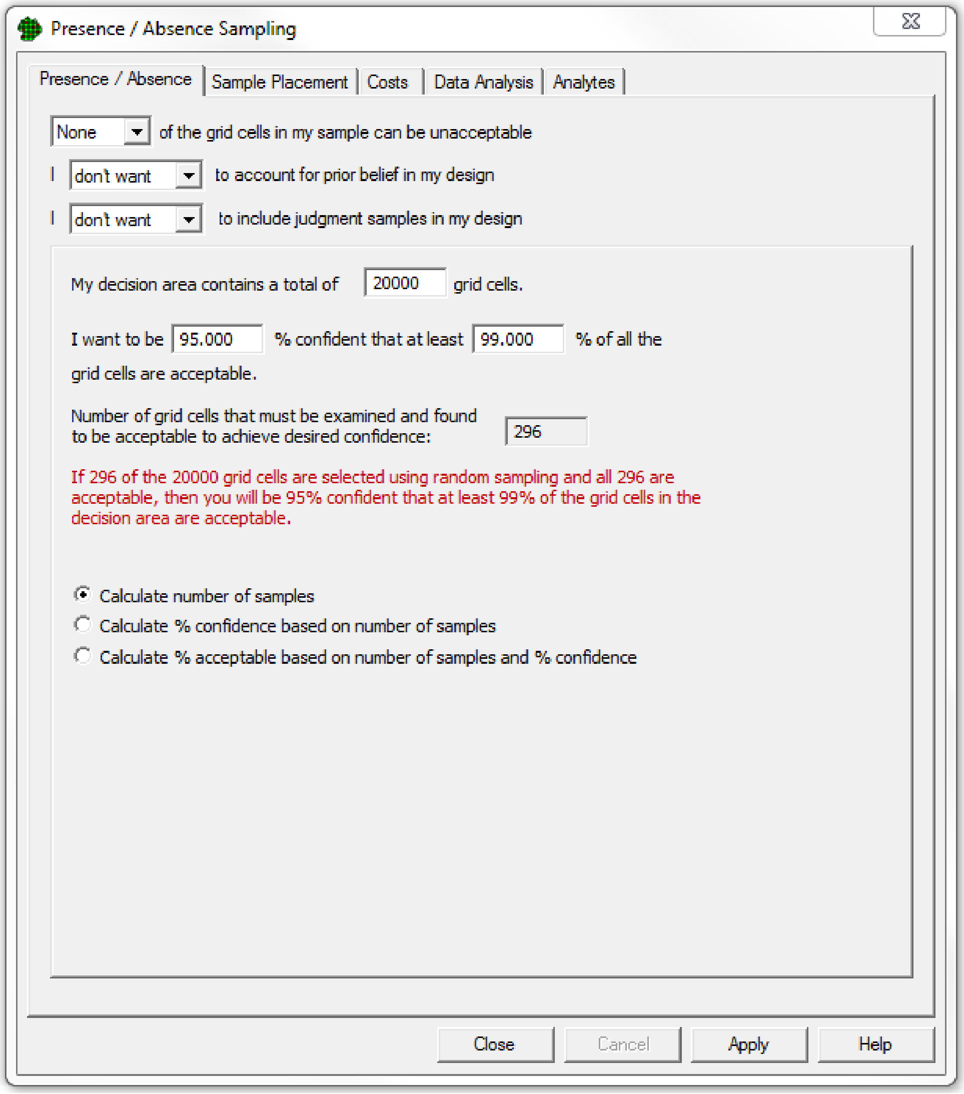

Figure 2.2.2. VSP screen giving the number of random samples needed when neither prior belief nor judgmental samples are included.

2.3 Clearance Sampling in VSP

The CJR method has five parameters that influence the number of required random samples. This section discusses how each influences the number of random samples and gives recommendations for reasonable values for each.

2.3.1 Recommendations for choosing *C* and *λ*

The value of the confidence level (C, λ) must ultimately be chosen by decision makers. It has been our experience that people usually feel more comfortable exploring a variety of options for the confidence level and then evaluating the corresponding tradeoffs with the random sample size.

The desired confidence (*C*) and the desired proportion of the decision area that is acceptable (λ) are usually determined in conjunction with each other. One approach would be to choose a value of *C* and then calculate the required sample sizes for various values of λ, and then choosing the largest value of λ that still yields a feasible sample size, provided $\lambda > C$. That is, we recommend a (C, λ) clearance statement be made about as large a percentage of the decision area being acceptable (i.e., the largest λ) as possible, provided the confidence level is suitably high. As an example, a (95%, 99%) clearance statement would be preferred to the much weaker (99%, 95%) clearance statement.

Another approach would be to fix the value of λ and then calculate the sample size for various values of *C*, where *C* is preferably less than λ. One may also approach the problem by fixing the sample size a priori and determining the (C, λ) combination that the sample size "buys." VSP facilitates all three of these approaches.

As an illustrative example, Figure 2.3.1 shows the number of random samples required for various levels of *C* and λ, given 20,000 grid cells, 30 judgmental samples , $P_J = 0.99$, and $r = 2$. As λ approaches 1, small increases in λ result in very large increases in the random sample size. In fact, among the six input parameters for the CJR model, the required number of random samples is the most sensitive to λ, which underscores the importance of identifying the smallest λ possible. While it would obviously be preferable to be able to state with high confidence that 100% of the population is acceptable, such large sample sizes are virtually never feasible in practice.

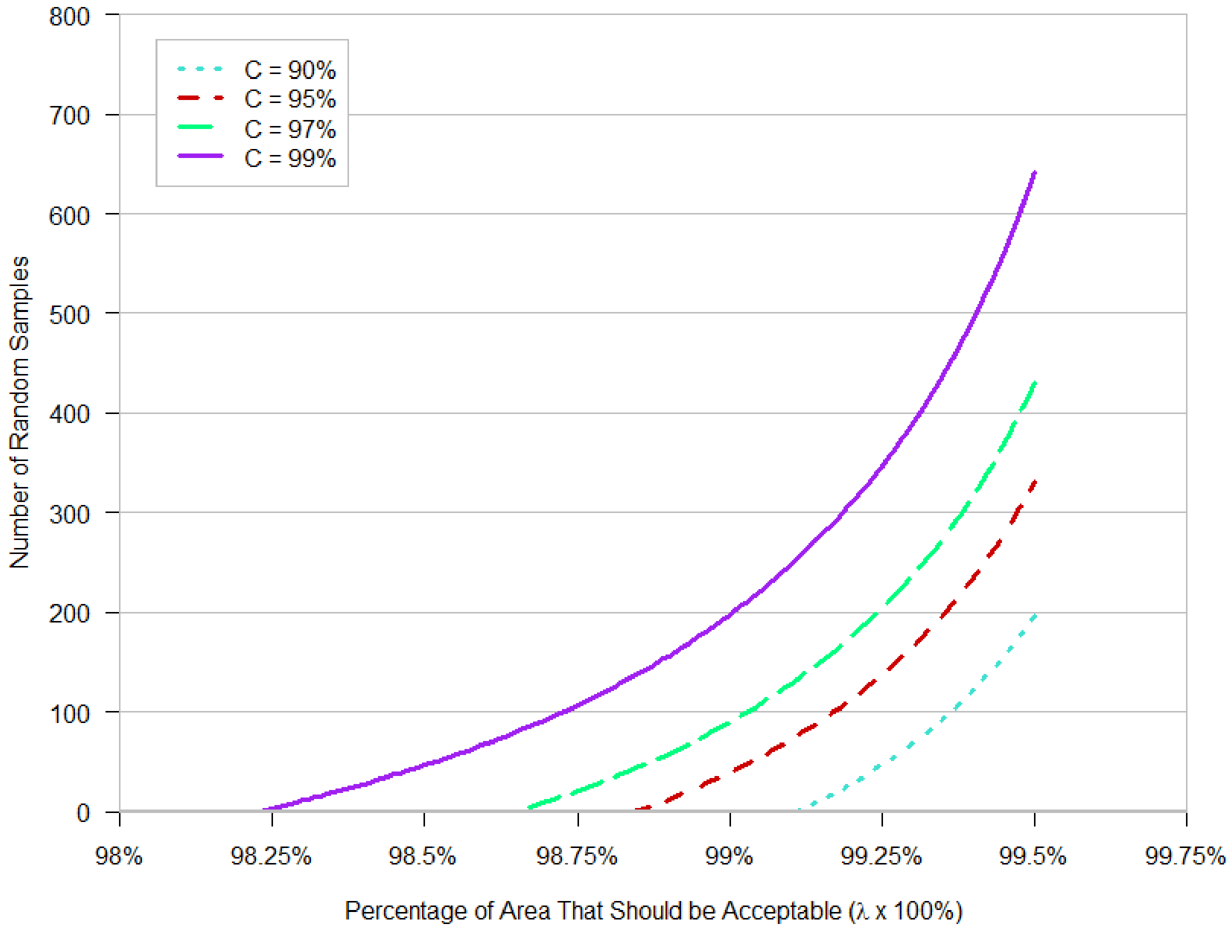

Figure 2.3.1. Required number of random samples versus λ for various values of C and for r = 2, P_J = 0.99, a population of size 20,000, and 30 judgmental samples.

2.3.2 Recommendations for P_J

The range of plausible values of P_J is 0.5 to 0.999. If the practitioner prefers to make a neutral (non-informative) assertion regarding the prior belief on P_J, setting P_J = 0.5 results in a uniform prior distribution. However, for acceptance sampling, the uniform prior is very conservative (and even pessimistic) because it essentially represents a prior belief where these two statements are equally likely:

A: the a priori chance of a high-risk grid cell being acceptable is 1%
B: the a priori chance of a high-risk grid cell being acceptable is 99%.

In the context of clearance sampling, statement B should be much more likely than statement A—else why would one even attempt clearance sampling? The notion of comparing the likelihoods' a priori chances may seem foreign to some. Admittedly, the idea takes some getting used to. It stems from the Bayesian model that underlies the CJR method, where the primary parameter of interest, the probability that a grid cell is acceptable, has its own (prior) probability distribution.

While is it obviously preferable to have some quantitative basis for choosing the value of P_J, perhaps based on previous experience or pilot studies, that information may not be available. However, the Bayesian construction of the CJR model does permit practitioners to rely on their subjective belief about the a priori acceptability of the decision area, based on what they know about the area. This prior knowledge may include the efficacy of a proven decontamination method, the likelihood a contaminant was present in the decision area, etc.

2.3.3 Recommendations for *r*

We recommend that r be chosen conservatively, so that practitioners err on the side of underestimating r. However, overly conservative estimates (when r is too small) can result in large numbers of random samples required to obtain the desired confidence level. In general, we recommend values of r between 1 and 5. Choosing a value of $r = 1$ makes the contribution of the judgmental samples equivalent to the random samples.

As an illustrative example, Figure 2.3.2 shows the number of random samples required for various levels of P_J and r, given 20,000 grid cells, 30 judgmental samples, $C = 0.95$ and $\lambda = 0.99$. This example shows a linear decrease in the number of random samples as r increases. It also shows that the random sample size decreases at a higher rate as P_J increases.

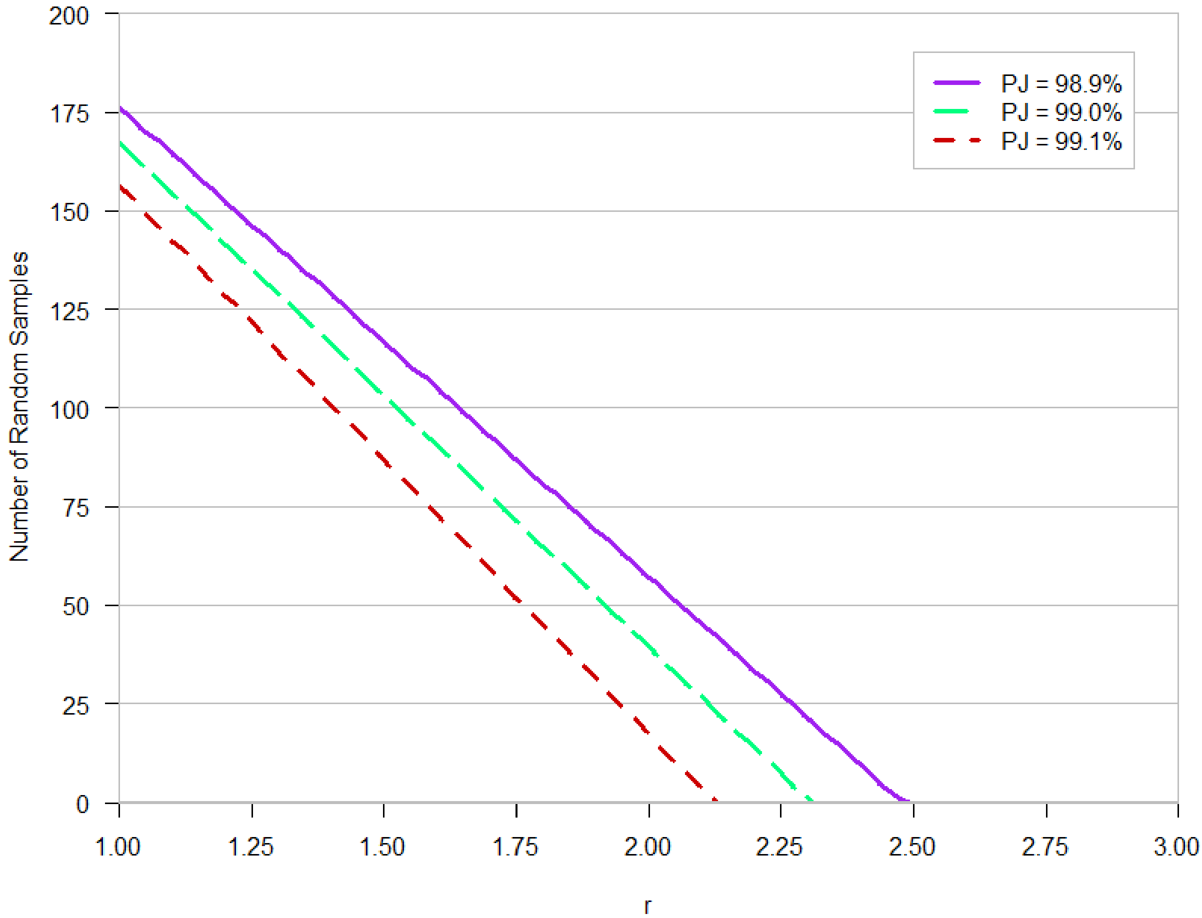

Figure 2.3.2. Random sample sizes versus *r* for various values of P_J and for C = 0.95, λ = 0.99, a population size of 20,000, and 30 judgmental samples.

2.3.4 Recommendations for the number of judgmental samples

The CJR method assumes that all high-risk grid cells are sampled—and the number of judgmental samples should be chosen to match the number of high-risk grid cells that are believed to exist in the decision area. While acceptable judgmental samples count for more in the CJR model than acceptable random samples, care should be taken to make sure that the high-risk grid cells are defined properly and not just for the purpose of minimizing the total number of required samples.

2.4 Model Assumptions

This section gives a summary of the assumptions required by the CJR methodology. In general, it is not feasible that every assumption be satisfied. However, more assumptions being satisfied yields greater confidence in the CJR methodology.

a. The total number of grid cells in the decision area is known and each grid cell is the same size.

b. The size of the grid cell is appropriate for the chosen sampling methodology. If more than one sampling methodology is employed in a decision area, the size of the grid cell is chosen to match the sampling methodology with the smallest footprint.

c. The outcome from each sample will be qualitative, i.e., acceptable or unacceptable.

d. Each sample is collected, handled, measured, or inspected using approved methods that yield sufficiently precise measurements.

e. The measurement (inspection) method correctly classifies each sample as being acceptable or unacceptable. That is, an acceptable grid cell is not classified as being unacceptable (a false positive) and an unacceptable grid cell is not classified as being acceptable (a false negative).

f. The potential measured outcomes of all grid cells are independent. If spatial correlation is present, the CJR method is conservative, (i.e., more samples are required than would otherwise be needed).

g. In the decision area, the high-risk grid cells are more likely to be unacceptable than the remaining grid cells which are low-risk grid cells.

h. All high-risk grid cells are sampled with judgmental samples.

i. A high-risk cell is, on average, r times more likely to be unacceptable than a low-risk cell.

j. Acceptable outcomes from judgmental samples increase the confidence that low-risk cells are also acceptable.

k. Before sampling takes place, the expected probability of a judgmental sample being acceptable is P_J.

l. The probability that a high-risk cell is unacceptable has a Beta prior distribution with the first shape parameter equal to 1 and the second shape parameter equal to $P_J / (1 - P_J)$. This ensures that the expected value of the Beta prior is $1 - P_J$.

m. A random sample is taken from the low-risk grid cells. The sample locations may be selected using simple random, systematic random, or adaptive fill sampling.

3.0 Special Considerations

There are certain combinations of the input parameters that practitioners should be aware of and generally avoid when using the CJR method, which we discuss below. We also discuss the implications of under or overestimating the number of high-risk cells, the prior belief, P_J, and the likelihood ratio between judgmental and random samples, r.

3.1 When P_J and/or r is large and λ is too low

When λ is small relative to the size of P_J or r,, the prior belief in the acceptability of the decision area may already satisfy the intent to demonstrate (through sampling) that λ x 100% of the decision area is acceptable. This can potentially result in fewer random samples being required in larger decision areas than in smaller decision areas. This phenomenon results from a characteristic of the Binomial cumulative distribution function that is used in the model, and most often occurs when P_J is close to 1 and/or r is large. To avoid this phenomenon, VSP identifies and recommends the *viable* value of the fraction of acceptable grid cells, λ_v, that is just large enough to ensure that, all things being equal, larger decision areas receive more samples than smaller decision areas. Practitioners can choose to use the recommended λ_v, but doing so may result in more samples than desired. Alternatively, they may prefer to reduce the values of P_J and/or r (which is conservative) which will decrease the value of λ_v. These smaller values of λ_v will result in smaller required sample sizes, albeit at the expense of lesser confidence.

3.2 "100% Acceptable" is not Acceptable

Although it is most desirable to state with high probability that the population has no unacceptable items (i.e. $\lambda = 1$), achieving that degree of certainty typically requires that almost all of the population cells be sampled. This is almost never feasible. Sego et al (2007) show that even with very strong prior values (i.e. large values of P_J), the number of random samples is still untenable. Figure 3.1 shows that reducing λ, if only a little, will produce a more acceptable number of random samples.

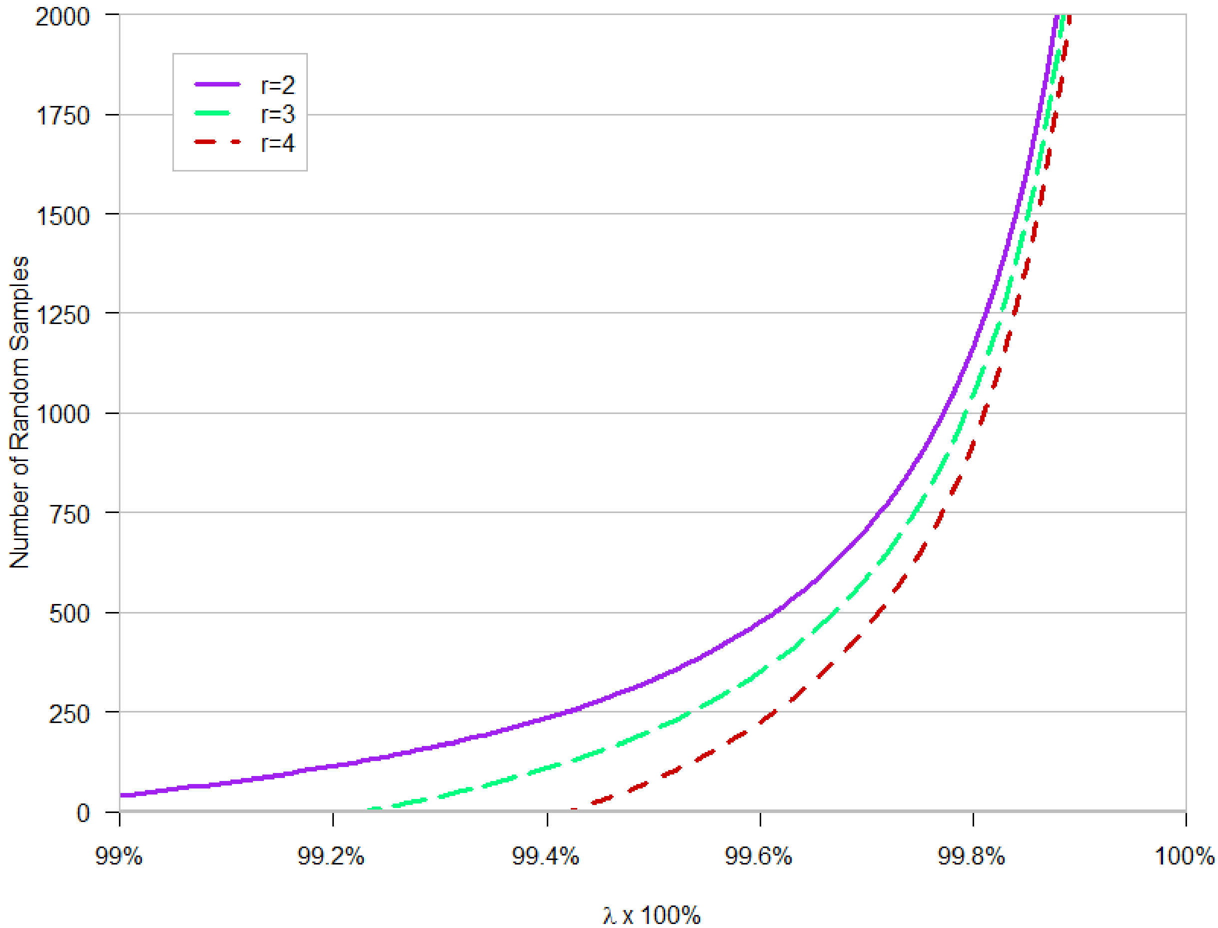

Figure 3.2.1. Random sample sizes versus λ for various values of r and for C = 0.95, P_J = 0.99, N = 20,000, and 30 judgmental samples.

3.3 Over or underestimating the number of high-risk cells

One of the assumptions of the CJR methodology is that all high-risk cells are sampled judgmentally. When the number of high-risk cells is overestimated, too many cells are sampled judgmentally and treated as high-risk cells. Provided all samples are acceptable, this gives practitioners an artificially high level of confidence that the decision area as a whole is acceptable. To the contrary, when the number of high-risk cells is underestimated, then not enough cells are sampled judgmentally, providing less evidence of "acceptability" than otherwise would be obtained—thus resulting in an artificially low confidence level. Hence, underestimating the number of high-risk cells is conservative.

3.4 Over or underestimating P_J

Similar to over or underestimating the number of high-risk cells, if P_J is overestimated, the CJR design will require fewer random samples than it should to achieve the desired level of confidence. Equivalently, the level of confidence will be artificially high for the number of samples. Conversely, if P_J is underestimated, the design will require more samples than necessary to achieve the desired confidence level. Equivalently, the level of confidence will be artificially low for the number of samples. Hence, underestimating P_J is conservative.

3.5 Over or underestimating r

Over and underestimating r has the same effect on the confidence and sample size as P_J and the number of high-risk cells. Overestimating r leads to artificially high confidence and too few samples. Underestimation gives rise to artificially low confidence and too many samples, and underestimating r is conservative.

3.6 Conclusion

As discussed in this paper, the CJR sampling methodology is a very flexible tool that can help practitioners defensibly sample an affected area, while taking as few random samples as possible. Augmenting high-risk judgmental samples with low-risk random samples allows for taking fewer total samples while maintaining a high level of confidence that the results from the sampling reflect the true acceptable state of a decision area.

4.0 References

Bowen MW and CA Bennett. 1988. *Statistical Methods for Nuclear Material Management*. NUREG/CR-4604. U.S. Nuclear Regulatory Commission, Washington, D.C.

Sego LH, KK Anderson, BD Matzke, WK Sieber, S Shulman, J Bennett, M Gillen, JE Wilson, and BA Pulsipher. 2007. *An Environmental Sampling Model for Combining Judgment and Randomly Placed Samples*. PNNL-16636, Pacific Northwest National Laboratory, Richland, Washington. Accessed March 8, 2012 at http://www.pnl.gov/main/publications/external/technical_reports/PNNL-16636.pdf.

Sego LH, S Shulman, KK Anderson, JE Wilson, BA Pulsipher, and WK Sieber. 2010. *Acceptance Sampling Using Judgmental and Randomly Selected Samples*. PNNL-19315, Pacific Northwest National Laboratory, Richland, Washington. Accessed March 8, 2012 at http://www.pnl.gov/main/publications/external/technical_reports/PNNL-19315.pdf.

Squeglia NL. 1994. *Zero Acceptance Number Sampling Plans*. ASQ Quality Press, Milwaukee, Wisconsin.

VSP Development Team. 2014. *Visual Sample Plan: A Tool for Design and Analysis of Environmental Sampling*, Version 6.6. Pacific Northwest National Laboratory, Richland, Washington. Available at http://vsp.pnnl.gov.

Pacific Northwest
NATIONAL LABORATORY

U.S. DEPARTMENT OF
ENERGY

Proudly Operated by Battelle *Since 1965*

902 Battelle Boulevard
P.O. Box 999
Richland, WA 99352
1-888-375-PNNL (7665)
www.pnnl.gov