

DESIGNING EXPERIMENTS THAT CONTROL FOR SPATIAL AND TEMPORAL VARIATION

PETER PETRAITIS

Department of Biology, University of Pennsylvania, Philadelphia, PA 19104-6018 USA

INTRODUCTION

Ecological data derived from experiments and observational studies done in the field are very often variable because patterns in nature are masked by temporal and spatial variation in physical and biological factors. Even though ecologists are well aware of the problems caused by natural variation, the ecological literature is filled with experiments that either ignore the difficulties caused by spatial and temporal variation or deal with them in an inappropriate manner. Here I provide an introduction to methods for controlling spatial and temporal variation. There is a vast literature in this area and I can only offer some guidelines and provide a list of helpful references.

Spatial and temporal variation can cause problems even in the most straightforward statistical tests. This can be easily seen in a simple example. Suppose we wished to test the hypothesis that grazing in the valleys that line Lake Hovsgol reduces aboveground biomass of herbaceous plants. We might test this hypothesis by setting up a simple *t*-test. The test is carried out in three steps. First, we pose the null hypothesis that the parametric means are equal (i.e. average aboveground biomass of plants is the same in grazed and ungrazed areas) and the alternative hypothesis that they are not. Next we randomly draw samples from the two populations - for example, areas where grazing animals have been excluded and areas where animals graze freely. Finally we test the hypotheses based on the magnitude of the test statistic *t*. The calculation of *t* is based on the sample sizes and sample estimates of the parametric means and variances. If we assume the two populations have the same variance and we draw samples of equal size, then $t = (\bar{Y}_1 - \bar{Y}_2)n^{1/2} / \sqrt{2}s$ where \bar{Y}_1 and \bar{Y}_2 equal the averages, *n* equals the sample size, and *s* equals the

standard deviation. Failure to consider spatial and temporal variation can cause us to mis-estimate the difference between the averages, $\bar{Y}_1 - \bar{Y}_2$, and/or the standard deviation, *s*.

Mis-estimation arises most often when the experiment is laid out incorrectly. For example in our grazing experiment, we should ideally have many replicate sites spread over a very large area and at various times of the year. In the half the sites, we should exclude grazing animals and in the other half animals should be allowed to graze freely. Making fences to keep out grazing animals is, however, quite expensive, and so for convenience, we might fence one large area. We then would take many "replicate" samples from within the single fenced area and compare these samples to a set of samples taken from an adjacent unfenced area. This sampling protocol, while easy and convenient, presents statistical problems because we may seriously mis-estimate the averages and the standard deviation (see Hurlbert 1984 for a complete discussion of this problem, which is known as pseudoreplication). The fenced area, for example, may be on a slightly lower and thus wetter site. Plant biomass in the fenced area would then be greater not only from reduced grazing but also from the effects of the added moisture, which tends to increase primary production. Thus the difference between the two averages confounds the effects of moisture and grazing pressure. In contrast, the standard deviation may be smaller than we expect because the samples within the fence area and in the adjacent area are so close together that they are correlated spatially. The end result is the *t*-test value will be too large and will be much larger than the test value if we had made many separate fenced areas and spread them over a large area. More importantly we may reject the null hypothesis but for the wrong reason. Our single fenced area may

differ from the single open area because of soil moisture and not grazing pressure. Yet spreading our sampling over a large area will most certainly increase the standard deviation and make it difficult to detect a difference over and above the background of environmental “noise.” Thus the problem is how to design an experiment that not only reflects the natural variation we see in the field but also controls for that variation so we are able to undertake statistical tests with sufficient power.

CONTROLLING FOR SPATIAL VARIATION

The effect of spatial variation is usually controlled by grouping observations and treatments into blocks. These blocks are groups of treatments that are placed nearby each other. In the simplest designs, blocks are transect lines laid across the area of interest with each line containing at least one replicate of each treatment. In analysis of variance terms, the design with one replicate of each treatment per block is known as a randomized block design, and designs with multiple replicates of each treatment per block are known as mixed models (for a fuller explanation, see Snedecor and Cochran 1989, Sokal and Rohlf 1995, Potvin 2001).

Transect lines or blocks can be laid out in a number of ways and our choice depends on our knowledge and preconceptions of how spatial variation affects physical and biological processes. In general, we usually know or assume replicate sites close together are likely to have similar responses because of similarities in local environmental conditions. Replicates within blocks (or on transect lines) are thus grouped together to standardize for the common local effects. The blocks themselves are spread apart so differences among blocks reflect differences among environmental conditions. If nothing is known about the pattern of spatial variation, transect lines or blocks are laid out at random. If we know something about the pattern of spatial variation, then the transects or blocks should be set out so replicate sites within each block or transect are under similar conditions. For example, if environmental conditions follow an elevational gradient, groups of replicate treatments can then be blocked by elevation. This is done by laying out transect lines that follow the elevational isoclines and randomly assigning the treatments (for example, fenced areas or open areas) to replicate

sites along the transect.

Three different analyses of a contrived example illustrates how blocking can control for spatial variation. Imagine we again wish to test the effects of grazing on plant biomass and we know that biomass varies with elevation. Figure 1A shows the contour map for grassland with 36 experimental sites spread across the elevational gradient. Grazing animals are excluded from half of the sites with fences and the other half is accessible to grazers. Assignment of treatments to sites should be done randomly, but to make the analyses of the examples clearer, the treatments are laid out in an alternating pattern. The contrived data are given in Table 1.

TABLE 1. Contrived data for aboveground plant biomass (kg/m^2). Layout of transects shown in Figure 1. Data were generated by summing elevation codes (I = 0.1, II = 0.2, etc.), treatment codes (Open areas = 0.1, Fenced areas = 0.2), and a random number between 0 and 0.1. For example, the entry for elevation V, transect B, and fenced = $0.5 + 0.2 + 0.0054 = 0.7054$.

	A	B	C	D	E	F
I	O 0.2626	F 0.3848	O 0.2866	F 0.3041	O 0.2449	F 0.3218
II	F 0.4128	O 0.3898	F 0.4627	O 0.3291	F 0.4251	O 0.3181
III	O 0.4851	F 0.5804	O 0.4290	F 0.5234	O 0.4722	F 0.5942
IV	F 0.6373	O 0.5191	F 0.6348	O 0.5774	F 0.6807	O 0.5827
V	O 0.6851	F 0.7054	O 0.6526	F 0.7886	O 0.6486	F 0.7700
VI	F 0.8867	O 0.7066	F 0.8883	O 0.7526	F 0.8778	O 0.7563

First suppose we know nothing about the potential effects of elevation on plant biomass. In this case, the 18 fenced sites and the 18 accessible sites would be considered replicate samples, and the simplest test would be a t-test or a one-way analysis of variance in which the spatial arrangement is ignored. Analysis of variance shows no significant effect of grazing because the error variance is too large (see Table 2). The average plant biomass between the two treatments appears to be different (Table 2) but there is too much “noise” (i.e. the standard errors are large) to detect a significant “signal” of the treatments (i.e. the difference between the treatments).

The outcome is quite different if information about elevation is included in the analysis. Suppose the transect lines follow the elevation contours exactly (Fig. 1B), and the analysis of

variance includes transect lines as blocks. The test now shows significant effect of grazing on plant biomass because the blocking has reduced the error variance (Table 2). Note the treatment averages for this analysis are identical to the averages from the first analysis (Table 2). In addition, the two analyses have the same treatment Sums of squares (Table 2, Treatment SS = 0.0881), which a measure of the amount of variation explained by the treatment effect. The “signal” remains the same but the level of noise has been reduced by taking into account the elevational gradient. As a result, the standard errors of the averages are reduced (Table 2). Note that the effect of blocking can be seen in the block averages (Fig. 2).

Blocking re-distributes the unexplained variation, which is related to the error Sum of Squares. In the one-way analysis of variance, in which spatial information is ignored, the error Sum of Squares is 1.1408 with 34 degrees of freedom and contains the variation due to elevation. The blocking by elevation in the second analysis of variance partitions the original error sums of squares into variation due to elevation (SS = 1.1127, d.f. = 5), to the interaction of treatment and elevation (SS = 0.0047, d.f. = 5), and to unexplained sources (i.e. the new error Sum of Squares, which equals 0.0234 with d.f. = 24). The sum of the new Sum of Squares and their degrees of freedom equal the original error Sum of Squares and its d.f. The variation is truly re-distributed by blocking.

Blocking can still improve the power of a test even when we don't know the pattern of spatial effects. As an extreme example, suppose we lay out the transects so they ran across the elevational gradient (Fig. 1C). Each block would then contain a collection of sites that differ quite bit but blocks will not differ very much from each other. There is no change in the treatment averages but the treatment standard errors are quite large (Table 2). The analysis of variance shows no significant difference between the treatments. The effect of blocks is also not significant, and the block averages are nearly identical (Fig. 2). The standard errors of the blocks are large because the spatial variation across the elevational gradient is contained within each block. This design does a poor job of accounting for the spatial effects of elevation, but it is a design that is commonly used when researchers do not understand the purpose of blocking, which controls for spatial variation by grouping similar sites together.

Normally we would expect our ability to control for spatial variation to fall somewhere the two extremes. At one extreme, the test is very powerful if we can account for the effects of spatial variation exactly. At the other extreme, the test is only slightly better than simple t-test if we completely misjudge the direction of the gradient. Even with a good guess of the direction of the gradient combined with random assignment of the direction and position of transects based on our guess will provide a reasonably powerful test.

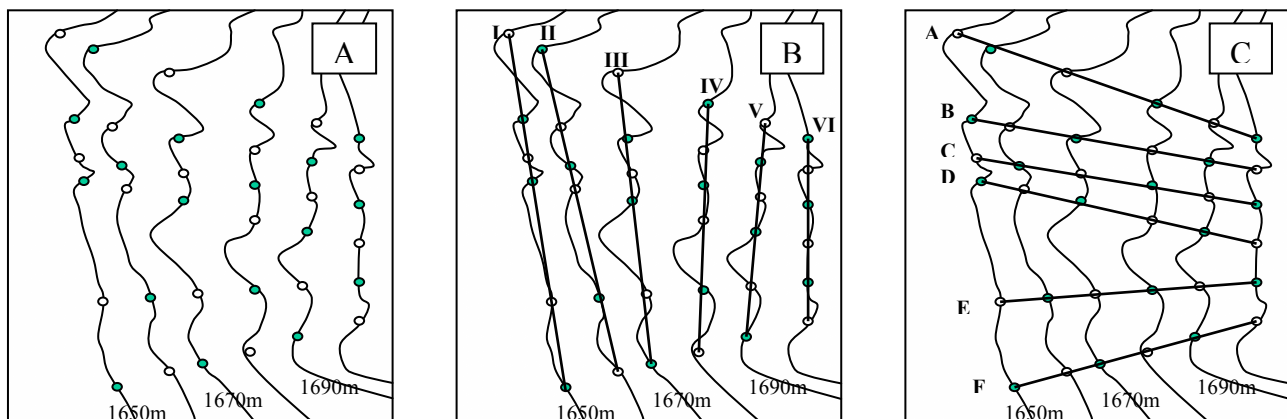


FIGURE 1. Contrived elevation contour maps showing positions of fenced and open areas. Shaded circles are fenced areas; open circles are open areas. Panel A shows layout for one-way analysis of variance that ignores spatial effects (Example 1 in Table 2). Panel B shows layout for two-way analysis of variance that controls for spatial variation by using elevation contours as blocks (Example 2 in Table 1). Panel C shows layout for two-way analysis of variance that fails to control for spatial variation because blocks (i.e. transects) run across the elevational gradient.

TABLE 2. Analyses of variances for examples. Data for examples 1, 2, and 3 are given in Table 1; for example 4 in Table 3. Layouts of experiments are found in the figures; examples 1, 2, and 3 in Figure 1, example 4 in Figure 3. d.f. = degrees of freedom, SS = Sums of Squares, MS = Mean Squares, P = probability levels. Numbers in parentheses in Example 4 are the adjusted degrees of freedom (Greenhouse-Geisser correction). Treatment averages (O = open areas; F = fenced areas) and standard errors (S.E.) are given with each analysis.

Source	d.f.	SS	MS	F-ratio	P	Average	S.E.
Example 1							
Treatment	1	0.0881	0.0881	2.625	0.1145	O 0.505	0.043
Error	34	1.1408	0.0336			F 0.604	0.043
Example 2							
Treatment	1	0.0881	0.0881	93.197	0.0002	O 0.505	0.007
Elevation	5	1.1127	0.2225	235.509	<0.0001	F 0.604	0.007
Treatment x Elevation	5	0.0047	0.0009	0.969	0.4562		
Error	24	0.0234	0.0010				
Example 3							
Treatment	1	0.0881	0.0881	5.342	0.0688	O 0.505	0.030
Transect	5	0.0013	0.0003	0.015	0.9998	F 0.604	0.030
Treatment x Transect	5	0.0824	0.0165	0.374	0.8613		
Error	24	1.0571	0.0440				
Example 4							
Between Subjects							
Analysis							
Treatment	1	0.0881	0.0881	662.08	<0.0001	O 0.505	0.003
Plot within Treatment =							
Error	4	0.0005	0.0001			F 0.604	0.003
Within Subjects Analysis							
Time (2.6)	5	1.1127	0.2225	194.657	<0.0001		
Treatment x Time (2.6)	5	0.0047	0.0009	0.8265	0.4935		
Error (10.6)	20	0.0229	0.0011				

CONTROLLING FOR TEMPORAL VARIATION

The control of temporal variation, at least naively, should be akin to the control of spatial variation with variation through time replacing variation across space. However, there are statistical difficulties because the data are often serially correlated. This means an observation from one time point is correlated with an observation from the next time point, and so the observations are not independent. For example, if the air temperature at 1300h is 10°C, it is likely the air temperature at 1400h will be within a couple of degrees of 10°C. The temperatures are correlated. Standard statistical tests, such as t-tests, will be biased if there are temporal correlations in the data because the sample variance is underestimated. The tests, thus, tend to reject the null hypothesis too

often. A full description of how to deal with serially correlated data is beyond the scope of this paper; see Box and Jenkins (1976) for a good introduction.

Repeated measures analysis of variance is often used to control for temporal variation (e.g., von Ende 2001). This is good method because the significance tests are adjusted for serial correlations. The adjustment, however, assumes a very specific form of serial correlation, and it can be difficult to assess how well the data meet the assumption about correlation structure. Repeated measures analysis of variance can be done as either a univariate or multivariate analysis, and these two approaches give different results because different assumptions are made about the correlation structure. It has been my experience the two approaches give similar results except when the

significance of a test is marginal.

The same contrived data can be used to illustrate how a repeated measures analysis of variance partitions the variation. Suppose we again wish to test the effects of grazing but we suspect the amount of plant biomass depends on the season. Thus early in the spring, we might expect plant biomass to be low in both open areas and fenced areas. As the season progresses and plants begin to grow, the amount of biomass should increase. However, we would expect the increase to be larger in the areas where grazers have been fenced out. Table 3 shows the contrived data from Table B re-arranged into a design of three fenced areas and three open areas, each sampled six times during the season.

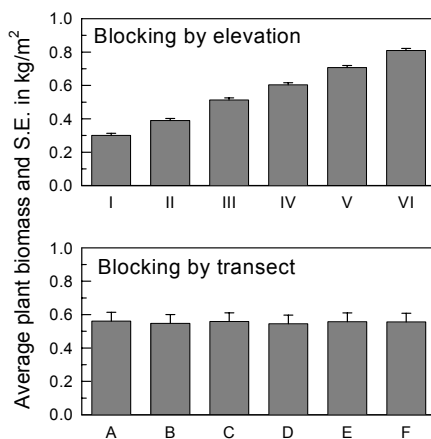


FIGURE 2. Block averages and standard errors from examples 2 and 3 in Table 2. Note block averages show the effects of elevation and have smaller standard errors when blocking is done by elevation.

Repeated measures analysis of variance divides the analysis into two parts - "Between Subjects" and "Within Subjects." This terminology comes from the field of experimental psychology, where repeated measures analysis was developed for the repeated testing of individuals (i.e., subjects). The Between Subjects analysis, tests the overall effect of the treatment and is akin to testing the averages across time. The Within Subjects analysis tests for trends across time and if that trend differs among treatments.

Table 2 gives the repeated measures analysis and shows the relationships to the standard analysis of variance in which elevation is blocked. Note that several of the Sums of Squares are identical because the same contrived data are used in both analyses. The treatment SS is the same in both analyses. Other Sums of Squares are the same but have different names: elevation SS = time SS and treatment x elevation SS = treatment x time SS. The error Sum of Squares in the two-way analysis of variance is partitioned into two parts in the repeated measures analysis of variance.

Although the Sums of Squares are similar, the interpretation of the tests differs slightly for the repeated measures analysis. The treatment F-ratio tests the null hypothesis that there is no overall average difference between fenced and open areas. The Between Subjects test is identical to an one-way analysis of variance using the averages across time as the replicates (i.e. 3 averages from the fenced areas and 3 averages from the open areas). Just as in the blocking of transects for spatial variation, controlling for the effects of time reduces the standard errors (Table 2).




FIGURE 3. Plot of contrived data used in the repeated-measures analysis. Shaded circles denote fenced areas; open circles denote open areas. Data are given in Table 3; analysis is given in Table 2.

The Within Subjects analysis tests for the effects of the repeated measurements done over time. The F-ratio for the effect of time tests the null hypothesis that there is no change in plant biomass over time. This test is significant, and we conclude that plants, not surprisingly, gain biomass as the

season progresses (see Fig. 3). The correction for serial correlation among the repeated measurement of the same plots is carried out by adjusting the degrees of freedom (see Table 2). Finally, the test for the effect of treatment x time asks if the average response in fenced and open areas differs over time. This test is not significant, suggesting that the temporal changes in plant biomass show a similar pattern. The only difference is that the fenced areas have higher yields at any point in time.

TABLE 3. Contrived data for repeated-measures analysis of variance. Plot of data is given in Figure 3. Data are identical to the data in Table 1, but re-arranged so that time replaces elevation. Codes for elevation, which are found in Table 1, are given below the new codes for time. Letters within the table identify block codes used in Table 1.

Treatment	Plot	Time 1	Time 2	Time 3	Time 4	Time 5	Time 6
		(I)	(II)	(III)	(IV)	(V)	(VI)
Open	1	0.2626 A	0.3898 B	0.4851 A	0.5191 B	0.6851 A	0.7066 B
Open	2	0.2866 C	0.3291 D	0.4290 C	0.5774 D	0.6526 C	0.7526 D
Open	3	0.2449 E	0.3181 F	0.4722 E	0.5827 F	0.6486 E	0.7563 F
Fenced	1	0.3848 A	0.4128 B	0.5804 A	0.6373 B	0.7054 A	0.8867 B
Fenced	2	0.3041 C	0.4627 D	0.5234 C	0.6348 D	0.7886 C	0.8883 D
Fenced	3	0.3218 E	0.4251 F	0.5942 E	0.6807 F	0.7700 E	0.8778 F

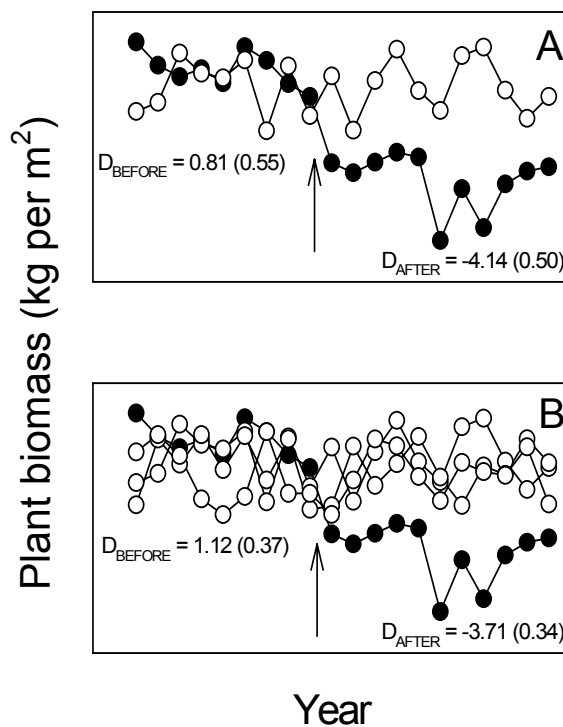


FIGURE 4. Contrived time series data with an intervention or change occurring at the time indicated by the arrows. Control sites, which are not affected by the intervention, are shown as open circles. The solid circles show the data for the single impact site. The D values are the averages based on the difference between the paired control and impact data at each time point. Numbers in parentheses are standard errors. The Before and After D's can be compared with a t-test. The t-test is often adjusted for the serial correlation in the data. Panel A shows standard BACI (Before-After-Control Impact) analysis, which is based on a single control site and a single impact site. Panel B shows a "Beyond BACI" design in which the D's are based on the difference between the average of several control sites and the single impact site at each time point.

MORE ON TEMPORAL VARIATION: INTERVENTION ANALYSIS, BACI AND BEYOND

One other situation is worth mentioned because it occurs so often in course of environmental monitoring. One could imagine a situation in which we have a single temporal series of observations before and after an event or change. For example we could have data on plant biomass before and after the introduction of grazing animals into a single area. The effect of grazing could be tested by comparing the before and after observations with a simple *t*-test. This kind of analysis is known as intervention analysis and was first used to test the effects of air pollution in Los Angeles after the implementation of laws for pollution controls (Box and Tiao 1975).

An extension of this approach involves not only data from before and after an intervention but also data from a control and an impact site (Fig. 4A). In the ecological literature, experimental designs are known as BACI designs (i.e., Before-After-Control-Impact designs). Analysis is carried out by taking the difference between paired control and impact observations and comparing the “before” differences against the “after” differences (see example in Fig. 4A).

In a series of papers, Underwood (1991, 1992, 1993, 1994) extended BACI designs to include multiple control and/or impact sites. These designs are known as “Beyond BACI” designs or IVRS design (for Impact Versus Reference Sites). The analysis is similar to the BACI analysis but the single control observation at each time point is replaced with the average of the reference (i.e. control) sites. The before and after differences are compared using a *t*-test or analysis of variance (see example in Fig 4B). Intervention analysis, BACI designs, and their extensions have subtle differences because of different assumptions about not only temporal variation but also spatial variation. For a full discussion of the issues raised by ecological data, see Underwood (1991, 1992, 1993, 1994), and Stewart-Oaten and Bence (2001).

RECOMMENDATIONS

In closing, I would like to offer several suggestions. While my advice is framed within the context of the problems raised by spatial and temporal variation, these are common-sense notions that are applicable to any field experiment.

1. Seek advice about experimental design and statistical analysis before you start an experiment, not after you collect the data.
2. Spread your replicates over the ranges of spatial and temporal variation that are of interest. Avoid setting up an experiment in which all the replicates of a particular treatment are grouped together in either time or space. For example, you should never use a single fenced area and repeatedly sampling plant biomass within that area.
3. Use your knowledge of spatial patterns to block groups of treatments together. Each block should encompass an area in which you expect the background environmental conditions to be similar. Each block should contain at least one replicate of each treatment.
4. Experimental designs with blocks are considered mixed models. F-ratios and associated degrees of freedom in mixed models depend on the number of blocks and treatments, not on the number of replicates per block by treatment combination. Given the choice between increasing the number of replicates or the number of blocks, you should always increase the number of blocks.
5. Data taken at different points in time are often serially correlated and analyses must correct for these correlations. Repeated measures analysis of variance is a reasonable approach if you wish to control for serial correlation.
6. Intervention analysis, BACI designs and their extensions can involve slightly different assumptions about spatial and temporal correlations. Consult someone familiar with these methods unless you are certain of how you wish to model spatial and temporal correlations.

USEFUL REFERENCES

Pseudoreplication

Hurlbert, S. H. 1984. Pseudoreplication and the design of ecological field experiments. *Ecological Monographs* 54:187-211.

Analysis of variance

Potvin, C. 2001. ANOVA: experimental layout and analysis. In: *Design and Analysis of Ecological*

- Experiments*, second edition, S. M. Scheiner & J. Gurevitch (ed.). Oxford University Press, Oxford, UK, pp. 63-76.
- Snedecor, G. W. & W. G. Cochran. 1989. *Statistical Methods*, eighth edition. The Iowa State University Press, Ames, Iowa, USA.
- Sokal, R. R. & F. J. Rohlf. 1995. *Biometry*, third edition. W.H. Freeman and Company, USA.
- Repeated measures analysis**
- Crowder, M. J. & D. J. Hand 1990. *Analysis of Repeated Measures*. Chapman & Hall, London, UK.
- Green, R. H. 1993. Application of repeated-measures designs in environmental impact and monitoring studies. *Australian Journal of Ecology* 18:81-98.
- von Ende, C. N. 2001. Repeated-measures analysis: growth and other time-dependent measures. In: *Design and Analysis of Ecological Experiments*, second edition, S. M. Scheiner & J. Gurevitch (ed.). Oxford University Press, Oxford, UK, pp. 134-157.
- Time series analysis**
- Box, G. E. P. & G. M. Jenkins 1976. *Time Series Analysis: Forecasting and Control*. Holden-Day, San Francisco, Calif. USA.
- Intervention analysis, BACI, and their extensions**
- Box, G. E. P. & G. C. Tiao 1975. Intervention analysis with applications to economic and environmental problems. *Journal of the American Statistical Association* 70: 70-79.
- Carpenter, S. R., T. M. Frost, D. Heisey & T. K. Kratz 1989. Randomized intervention analysis and the interpretation of whole-ecosystem experiments. *Ecology* 70:1142-1152.
- Ellis, J. I. & D. C. Schneider. 1997. Evaluation of a gradient sampling design for environmental impact assessment. *Environmental Monitoring and Assessment* 48:157-172.
- Rasmussen, P. W., D. M. Heisey, E. V. Nordheim & T. M. Frost. 2001. Time series intervention analysis: unreplicated large-scale experiments. In: *Design and Analysis of Ecological Experiments*, second edition, S. M. Scheiner & J. Gurevitch (ed.). Oxford University Press, Oxford, UK, pp. 158-177.
- Stewart-Oaten, A. & J.R. Bence 2001. Temporal and spatial variation in environmental impact assessment. *Ecological Monographs* 71: 305-339.
- Stewart-Oaten, A., J. Bence & C. Osenberg. 1992. Assessing effects of unreplicated perturbations: no simple solutions. *Ecology* 73:1396-1404.
- Stewart-Oaten, A. W. Murdoch & K. R. Parker. 1986. Environmental impact assessment: "pseudoreplication" in time? *Ecology* 67:929-940.
- Underwood, A. J. 1991. Beyond BACI: experimental designs for detecting human environmental impacts on temporal variation in natural populations. *Australian Journal of Marine and Freshwater Research* 42:567-587.
- Underwood, A. J. 1992. Beyond BACI: the detection of environmental impacts on populations in the real, but variable, world. *Journal of Experimental Marine Biology and Ecology*. 161:145-178.
- Underwood, A. J. 1993. The mechanics of spatially replicated sampling programmes to detect environmental impacts in a variable world. *Australian Journal of Ecology* 18:99-116.
- Underwood, A. J. 1994. Beyond BACI: sampling designs that might reliably detect environmental disturbances. *Ecological Applications* 4:3-15.