DETECTING ECOLOGICAL IMPACTS

Concepts and Applications in Coastal Habitats

Edited by

Russell J. Schmitt
Department of Biological Sciences
University of California
Santa Barbara, California

Craig W. Osenberg
Department of Zoology
University of Florida
Gainesville, Florida

Academic Press
San Diego New York Boston
London Sydney Tokyo Toronto

CHAPTER 1

DETECTING ECOLOGICAL IMPACTS CAUSED BY HUMAN ACTIVITIES

Craig W. Osenberg and Russell J. Schmitt

Ecologists and environmental scientists have long sought to provide accurate scientific assessments of the environmental ramifications of human activities. Despite this effort, there remains considerable uncertainty about the environmental consequences of many human-induced impacts, particularly in marine habitats (e.g., NRC 1990, 1992). This is especially surprising when one considers the vast amounts of capital and human resources that have been expended by industry, government, and academia in reviewing, debating, and complying with, a plethora of environmental regulations, which often require extensive study and documentation of environmental impacts. As we face an ever increasing number of environmental problems stemming from human population growth, it is critical that we achieve better understanding of the effects of humans. Due to the variety of human activities that potentially affect ecological systems, it also is imperative that we discriminate among effects of specific types of disturbances (rather than focus on an overall effect without regard to the particular sources), so that we can identify and give adequate attention to those that are most "harmful" (e.g., having the biggest effects, or which affect the most "valuable" resources). This requires approaches that can isolate effects of particular activities from nonhuman sources of natural variation as well as background variation caused by other anthropogenic events. Such approaches should reduce the uncertainty that underlies the documentation of effects of anthropogenic impacts and thus facilitate solutions to many of these problems.

Uncertainty surrounding the effects of anthropogenic activities arises from limitations imposed during the two scientific processes that comprise environmental impact assessment: (i) the predictive process, aimed at detailing the likely impacts that would arise from a proposed activity (most recently termed "Risk Assessment"; Suter 1993), and (ii) the postdictive process, aimed at quantifying the actual impacts of an activity (sometimes called "retrospective risk

Detecting Ecological Impacts: Concepts and Applications in Coastal Habitats, edited by R. J. Schmitt and C. W. Osenberg

assessment," and which we will refer to as "Field Assessment"). Instead of standing alone, these two processes should proceed in tandem and build upon each other; resolution of many environmental issues requires both quantification of impacts, as well as accurate prediction of future impacts. Neither process substitutes for the other. For example, a prediction reveals little unless the prediction is an accurate indicator of actual effects; the accuracy of predictions can only be assessed after repeated tests (i.e., comparison with the actual outcomes). Similarly, documenting an impact that has already occurred yields only a limited ability to improve environmental planning (i.e., by avoiding environmental problems, or facilitating environmentally safe activities) unless we use this information to construct or refine frameworks that enable us to accurately anticipate future environmental impacts (e.g., predict their magnitude, and know the likelihood of such impacts based on the type of activity, its location, or the system being affected). The development of such frameworks is crucial to sound decisions being made before an activity occurs.

To date, we have only a limited ability to accurately predict the ecological consequences of many anthropogenic impacts (e.g., Culhane 1987, Tomlinson and Atkinson 1987, Buckley 1991a, 1991b, Ambrose et al., Chapter 18). For example, audits of environmental impact assessments (i.e., comparisons of predicted impacts with those actually observed) often have found relatively good agreement between predictions and reality when focused on physical or engineering considerations (e.g., the amount of copper discharged from a wastewater facility), but poor or limited agreement when focused on biological considerations (e.g., impacts on population density). As Ambrose et al. (Chapter 18) point out, the poor agreement stems both from lack of quantitative (or often qualitative) predictions of ecological change in the predictive phase of assessments, as well as lack of knowledge of the actual impacts (due to the absence or poor design of Field Assessments). Advances in Risk Assessment are likely to promote more precise predictions and thus reduce the first hurdle; however, the second will continue to plague us until Field Assessment studies are designed that better isolate effects of human activities from other sources of variation. Furthering this latter goal is a major theme of this book.

The field of environmental impact assessment is quite broad, requiring expertise from a diversity of fields, including physics, chemistry, engineering, toxicology, ecology, sociology, economics, and political science. In assembling the components of this book, we did not seek a comprehensive treatise on impact assessment. Instead, we focused on a narrower but central topic: the Field Assessment of localized impacts that potentially affect ecological systems (we further use marine habitats to provide the context for the discussions). We chose this conceptual focus because a wealth of books have appeared in the past 10 years that deal well with other aspects of environmental assessment (e.g., Risk Assessment: Bartell et al. 1992, and Suter 1993; general overviews of EIA: Westman 1985, Wathern 1988, Erickson 1994, Gilpin 1994; general introduction to monitoring: Spellerberg 1991; see also Petts and Eduljee 1994). None of these books, however, deals with the issue of Field Assessments in more than a cursory way (but see NRC 1990 for a good introduction to the role of environmental monitoring); there are no detailed discussions of sampling designs that can most reliably estimate the magnitude of the impacts and quantify the power of the designs to detect impacts, nor are there evaluations of institutional and scientific constraints that limit the application of such designs. This book is designed to fill that gap, and in so doing, provides grist for future discussion and advances that are critically needed to better understand effects of anthropogenic activities.

Two major tenets, which we elaborate below, underlie this book: (i) that Field Assessments are absolutely essential to understanding human impacts, in part, because they complement, and provide field tests of, predictions provided by Risk Assessment; and (ii) that improved sampling designs are critical to improving the quality and utility of results obtained from Field Assessments.

The Need for Field Assessments

The emerging field of Ecological Risk Assessment (Bartell et al. 1992, Suter 1993) has led to a tremendous increase in the precision and explicitness of predictions of anthropogenic impacts on ecological systems. These predictions are often based on models derived from laboratory studies of toxicological effects, transport models that describe the movement of contaminants, and population models that attempt to couple physiological and demographic changes with shifts in population dynamics and abundances. However, no degree of sophistication of such models can guarantee the accuracy of the predictions. The quality and applicability of Risk Assessment can only be judged by the degree to which its predictions match the impacts that actually occur. This requires estimation of the magnitude of the impact, not just its detection, and thus requires a Field Assessment that is able to separate natural spatial and temporal variability from variation imposed by the activity of interest. This is not a trivial problem, and many previous assessments have failed in this regard. The small number of successes that exist are too few to permit any sort of rigorous evaluation of Risk Assessment models.

As Risk Assessment models become more complex and sophisticated, it is possible that they will be championed as the final step in environmental impact assessment; follow-up Field Assessments might be deemed a waste of effort (redundant with the effort devoted to obtaining the predictions). While this is a worthy (but elusive) goal, no Risk Assessment model, no matter how sophisticated, is currently capable of accurately predicting ecological change in response to an anthropogenic activity. As mentioned above, this, in part, is due to the lack of knowledge about the actual response of many systems to anthropogenic disturbances, and therefore the general inability to compare predicted and observed change.

Uncertainty in predictions from Risk Assessment models often is acknowledged but typically is limited to two sources: (i) uncertainty about the actual value of parameters that are estimated from the studies that underlie the Risk Assessment model; and (ii) uncertainty about the environmental inputs to the model (e.g., how much freshwater runoff will enter an estuary during an upcoming year). Estimates of these sources of error often are incorporated into a Risk Assessment analysis to estimate the uncertainty associated with the prediction(s) of the model. Another, perhaps more important source of uncertainty rarely is examined: the uncertainty that the model chosen exhibits dynamics that are quantitatively (or even qualitatively) similar to the dynamics exhibited by the actual system. For example, improved laboratory techniques might provide improved quantification of the effect of a toxicant on the fecundity of a focal organism. This information might then be used in a model that links toxicant exposure with fecundity, and fecundity with population growth. However, even if the effect of the toxicant can be accurately extrapolated to field conditions, there is little guarantee that the connection between fecundity and population dynamics has been modeled correctly. More generally, the predicted dynamics may bear little resemblance to the observed dynamics, not because of uncertainty in the laboratory measurements, but due to uncertainty in the structure of the model into which the laboratory data are embedded. Addressing this uncertainty requires extensive field data, including information on the link between physiological changes and behavior (e.g., habitat selection, mate selection, reproductive condition), demographic consequences (e.g., changes in survival, birth rates, migration), population-level responses (e.g., shifts in age-structure, temporal dynamics), community responses (e.g., due to shifts in the strengths of species interactions), and ecosystem properties (e.g., feedbacks between biotic shifts and the physio-chemical aspects of the environment). Ultimately, these field data, together with laboratory data and the Risk Assessment model(s), must be integrated and then tested via comparison with actual responses to specific human activities. This last step requires properly crafted Field Assessment designs of sufficient power to distinguish the effects of the activity from a diverse set of other processes that drive variation in ecological systems.

The (In-)Adequacy of Exisitng Field Assessment Designs

The Goal of Field Assessments

A basic goal of a Field Assessment study is to compare the state of a natural system in the presence of the activity with the state it would have assumed had that activity never occurred. Obviously, we can never know, or directly observe, the characteristics of a particular system (occupying a specific locale at a specific time) in both the presence and absence of an activity. Thus, fundamental goals of the assessment study are to estimate the state of the system that would have existed had the activity not occurred, estimate the state of the system that exists with the activity, and estimate the uncertainty associated with the difference

between these estimates (Stewart-Oaten, Chapters 2 and 7). The inability of most studies to accomplish these goals has, in large part, led to tremendous uncertainty regarding the environmental consequences of anthropogenic activities. We briefly review some of these design considerations, beginning with an often misunderstood approach borrowed from modern field ecology—the manipulative field experiment.

The Role of Field Experiments

Manipulative field experiments (with spatial replication of independent subjects, and randomized assignment of subjects to treatment groups) is a common and powerful tool of field ecologists. However, field experiments can do very little to resolve the specific goal of Field Assessments. This issue (i.e., the application of experimental design to assessment) has clouded much of the debate about the design of Field Assessment studies (e.g., Hurlbert 1984, Stewart-Oaten et al. 1986). While field experiments may provide crucial insight into the functioning of systems and the role of particular processes (typically acting over limited spatial and temporal scales), a field experiment cannot reveal the effects of a specific activity on the system at a specific locale at a specific time, which is often the focus of a Field Assessment. A field experiment could provide a powerful way to determine the average effect of a process (e.g., an anthropogenic activity) defined over replicates drawn at random from a larger population of study (assuming that we could conduct such a replicated experiment on the appropriate spatial and temporal scale). However, this field experiment could not tell us about the effect of the treatment on any one of the replicates. Yet, this is analogous to the problem faced in Field Assessments.

To illustrate, consider the possible environmental impacts related to offshore gas and oil exploration, specifically those associated with the discharge of drilling muds. We could conduct an experiment to address whether oil exploration has localized effects on benthic infauna inhabiting a particular region, say the Southern California Bight, by (i) randomly selecting a subset of sites within the Bight and allocating these between "Control" and "Treatment" groups; (ii) drilling exploratory wells and releasing muds in our Treatment sites; (iii) quantifying the abundances of infauna in the Control and Treatment sites after a specified amount of time (e.g., 1 year); and (iv) determining if there is sufficient evidence to reject the null hypothesis of "no effect" (e.g., are the means of the two groups sufficiently different to be unlikely to have arisen by chance?) using standard statistical procedures (e.g., a *t*-test).

Clearly, this is an unlikely scenario (few oil companies would be willing to have a group of ecologists dictate where they will conduct their exploration), but in some situations, such an opportunity might exist. If so, then we will be in a tremendous (and enviable) position to estimate the average local effect of oil exploration on benthic fauna inhabiting the Southern California Bight. While such a study would provide invaluable information, the results would say nothing about the effect of a single oil platform drilling exploratory wells at a specific site within the Bight. Indeed, it is possible that a significant (and biologically important) treatment effect could be found in our experiment even if there were no effects at a majority of the Treatment sites (so long as the remaining Treatment sites were sufficiently affected). There are, of course, situations where knowledge of the average affect of an activity would be quite useful (e.g., the administrative process of Environmental Impact Assessment). However, in most Field Assessments we are less concerned with the average effect and more concerned with the specific effects of a particular project at a specific locale. This is analogous to the experimentalist pondering the effect of the treatment on a single replicate (and not a collection of replicates).

Furthermore, an oil company certainly does not randomly select sites for exploration. It always could be argued that part of the selection criteria included the need to find sites that not only yield oil or gas, but also are sites in which oil and gas could be found and extracted without any environmental damage; resolution of the issue thus requires specific information about specific locales. Therefore, we require a tool more powerful, or at least more specific, than the replicated field experiment with randomized assignment.

Instead of manipulative field experiments, the basic tools used in Field Assessments involve monitoring of environmental conditions. Many such monitoring designs bear superficial resemblance to one another, but differ in some fundamental aspects. In the next section, we draw on discussions from Underwood (1991) and Osenberg et al. (1992) to clarify some of these distinctions. We illustrate the basic elements of the most commonly used assessment designs, and summarize results from studies with which we have been associated to illustrate where these studies can go wrong.

The Control-Impact Design

Perhaps the most common Field Assessment design involves the comparison of a Control site (a place far enough from the activity to be relatively unaffected by it) and an Impact site (i.e., near the activity and thus expected to show signs of an effect if one exists); a common variant involves a series of Impact sites that vary in their proximity to the activity. This sort of design often is part of the monitoring program required by regulatory agencies for various coastal activities. Environmental parameters typically are sampled at the two sites (with multiple samples taken from each site), and an "impact" is assessed by statistically comparing the parameters at the Impact and Control sites.

We illustrate this approach in Figure 1.1a, which shows that the density of a large gastropod (*Kelletia kelletii*) was significantly greater at a Control site (1.6 km from a wastewater diffuser) than at either of two Impact sites (located 50 and 250 m from the diffuser). This difference might be taken as evidence that the discharge of wastewater had a negative effect on the density of the gastropod.

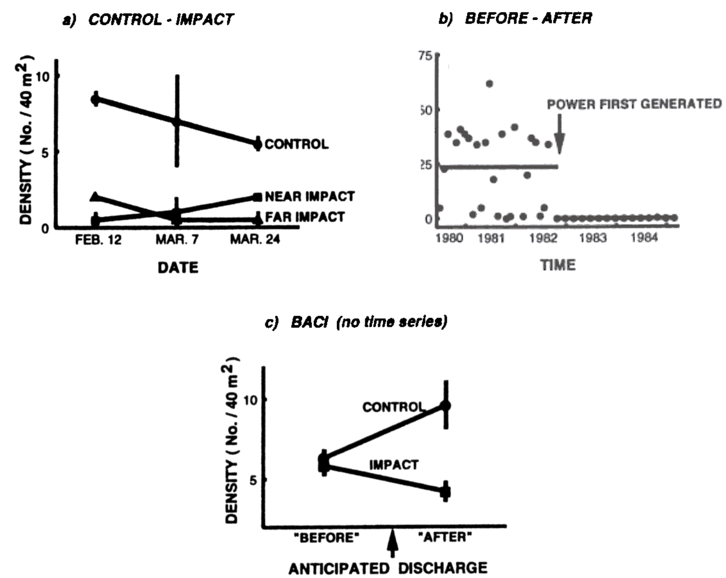

Figure 1.1. Three commonly used assessment designs that confound natural variability with effects of the anthropogenic activity. (a) The Control-After design showing the density of the snail *Kelletia kelletii* at three sites over time. The Near (square) and Far (triangle) Impact sites are located 50 and 250 m downcurrent of a wastewater outfall; the Control site (circle) is 1500 m upcurrent. These data were collected prior to discharge of wastewater. Shown for each date are the mean and range of gastropod density (n = 2 band transects per site). (b) The Before-After design showing density (catch per otter trawl) of pink surfperch *Zalembius rosaceus* over time at a location 18 km from the San Onofre Nuclear Generating Station (SONGS). The arrow indicates the first date on which power was generated by two new units of SONGS. Mean densities during the Before and After periods are indicated by the solid lines. (c) The Before-After-Control-Impact (BACI) design of Green (1979) showing the density of seapens *Acanthoptilum* sp. at two sites. The Control site is located 1500 m upcurrent, and the Impact site 50 m downcurrent, of a wastewater outfall. Because of permitting and production delays, discharge of wastewater did not begin when expected; all data were collected prior to discharge. Shown are means (± SE) using all observations within a period as replicates. The figure is adapted from Osenberg et al. (1992).

However, these data were collected prior to the discharge of wastewater. Thus, these differences observed during the Before period simply indicate spatial variation arising from factors independent of the effects of wastewater. To be applied with confidence, the Control-Impact design requires the stringent and unrealistic assumption that the two sites be identical in the absence of the activity. However, ecological systems exhibit considerable spatial variability, and it is extremely unlikely that any two sites would yield exactly the same result if sampled sufficiently. This design fails to separate natural *spatial* variability from effects of the activity.

The Before-After Design

An alternative design requires sampling of an Impact site both Before and After the activity; this avoids problems caused by natural spatial variation. Here, a significant change in an environmental parameter (e.g., assessed either by comparison of one time Before and one time After using within site sampling error as a measure of variability, or sampled several times Before and After and using the variation in parameter values through time as the error term) is taken as evidence of an "impact". Figure 1.1b provides an example for a fish, pink surfperch (*Zalembius rosaceus*), sampled Before and After the generation of power from new, seawater-cooled units of a large nuclear power plant (DeMartini 1987). The precipitous decline in abundance of pink surfperch is suggestive of a dramatic and detrimental impact from the power plant. However, these data are from a Control site 18 km from the power plant (a similar pattern also was seen at an Impact site: DeMartini 1987). Instead of indicating an impact, these data simply reflect the effect of other processes that produce temporal variability (in this case, it was an El Niño Southern Oscillation event that began at the same time as initiation of power generation: Kastendiek and Parker 1988). Applied in this way, the Before-After design fails to separate natural sources of *temporal* variability from effects of the activity.

A more sophisticated Before-After design is possible, and a classic example of intervention analysis (Box and Tiao 1975) provides both an illustration of its successful application and helps identify why the approach is limited for most ecological studies. Box and Tiao estimated the influence of two interventions (a traffic diversion and new legislation) on the concentration of ozone in downtown Los Angeles. Their procedure required that they (i) frame a model for the expected change; (ii) determine the appropriate data analysis based on this model; (iii) diagnose the adequacy of the model and modify the model until deficiencies were resolved; (iv) make appropriate inferences. Their analysis provided estimates of the effect of each intervention on ozone concentration.

There are several features of their system/problem that facilitated their successful analysis: (i) there was a long and intensive time series of ozone samples (hourly readings were available over a 17-year period, which included several years during the pre- and postintervention periods); (ii) the dynamics of ozone concentration were fairly well behaved, with repeatable seasonal and annual patterns; (iii) the number of pathways for the production and destruction of ozone were relatively few. These features contrast markedly with many ecological systems, where (i) we often have little expectation of how the system is likely to respond; (ii) data are sparse (time series are short, and intervals between sampling are long); and (iii) population density (for example) can be influenced by a multitude of processes (including a variety of mechanisms driven by abiotic factors and a wealth of mechanisms involving interactions with other species, each of which is also influenced by a variety of factors, including the effects of the focal activity). Certainly, such an approach might provide a powerful way to assess responses of biological systems to interventions (Carpenter 1990, Jassby and Powell 1990), but currently it remains limited due to the paucity of detailed knowledge about dynamics of ecological systems. In cases where data from an unaffected Control site are available, we may be able to incorporate them into such time series analyses to compensate for the sparseness and complexity of ecological data (Stewart-Oaten, Chapter 7: see below).

Before-After-Control-Impact (BACI) Designs

One possible solution to the problems with the Control-Impact and Before-After designs is to combine them into a single design that simultaneously attempts to separate the effect of the activity from other sources of spatial and temporal variability. There are a variety of such designs. In the first, which we refer to simply as BACI (Before-After-Control-Impact), a Control site and an Impact site are sampled one time Before and one time After the activity (Green 1979). The test of an impact looks for an interaction between Time and Location effects, using variability among samples taken within a site (on a single date) as the error term. Data from our studies of a wastewater outfall (Figure 1.1c) demonstrate such an interaction; the decline in the density of seapens (*Acanthoptilum* sp.) at the Impact site relative to the Control site suggests that the wastewater had a negative effect on density of seapens. However, discharge of wastewater at this site was delayed several years, and did not occur when first anticipated. Thus, the observed changes were due to other sources of variability and were not effects of the wastewater. This design confounds effects of the impact with other types of unique fluctuations that occur at one site but not at the other (i.e., Time × Location interactions). Unless the two sites track one another perfectly through time, this design will yield erroneous indications that an impact has occurred.

To circumvent this limitation of Green's BACI design, Stewart-Oaten et al. (1986; see also Campbell and Stanley 1966, Eberhardt 1976, Skalski and McKenzie 1982) proposed a design based on a time series of differences between the Control and Impact sites that could be compared Before and After the activity begins. We refer to this design as the Before-After-Control-Impact Paired Series (BACIPS) design to highlight the added feature of this scheme (see Stewart-Oaten, Chapter 7 and Bence et al., Chapter 8). In the original derivation of this design (e.g., Stewart-Oaten et al. 1986), the test of an impact rested on a comparison of the Before differences with the After differences. Each difference in the Before period is assumed to provide an independent estimate of the underlying spatial variation between the two sites in the absence of an impact. Thus, the mean Before difference added to the average state of the Control site in the After period yields an estimate of the expected state of the Impact site in the absence of an impact during the After period: i.e., the null hypothesis. If there were no impact, the mean difference in the Before and After periods should be the "same" (ignoring sampling error). The difference between the Before and After differences thus provides an estimate of the magnitude of the environmental impact (and the variability in the time series of differences can be used to obtain confidence intervals: Stewart-Oaten, Chapter 7 and Bence et al., Chapter 8).

The BACIPS design is not without its limitations, for it also makes a set of assumptions, which if violated can lead to erroneous interpretations (e.g., due to nonadditivity of Time and Location effects or serial correlation in the time series of differences). Indeed, one of the fundamental contributions of Stewart-Oaten's work (Stewart-Oaten et al. 1986, 1992, Stewart-Oaten, Chapter 7)has been to make explicit the assumptions that underlie the BACIPS design, pointing out the importance of using the Before period to generate and test models of the behavior of the Control and Impact sites, and to suggest possible solutions if some of the assumptions are violated. Importantly, many of these assumptions can be directly tested. Of course, it is still possible that a natural source of Time × Location interaction may operate on the same time scale as the study, and thus confound interpretation of an impact. However, this problem is far less likely than those inherent to the other designs (e.g., that variation among Times and Locations be absent and that there be no Time × Location interaction).

In this volume, Stewart-Oaten (Chapter 7) and Bence et al. (Chapter 8) elaborate upon and apply a more flexible BACIPS design based on the use of the Control site as a "covariate" or predictor of the Impact state, which might have even greater applicability than the original design (which was based on the "constancy" of the differences in the Before and After periods). Underwood (1991, Chapter 9) has suggested a "beyond-BACI" approach, which incorporates multiple Controls, as well as random sampling of the study sites (thus, the "Paired Series" aspect of the BACIPS design is not present in Underwood's beyond-BACI design). Underwood suggests that the beyond-BACI design is able to detect a greater variety of impacts than the BACIPS design (e.g., detection of pulse responses as well as sustained perturbations); however, he also notes that his design is not able to deal explicitly with problems of serial correlation. By contrast, the presence of serial correlation can be directly assessed, and appropriate action taken, when applying the BACIPS design (Stewart-Oaten et al. 1986, 1992, Stewart-Oaten, Chapter 7). In a variety of important ways, Underwood's approach differs from Stewart-Oaten's and some others represented in this book (e.g., Osenberg et al., Chapter 6, Bence et al., Chapter 8). While both schools-of-thought advocate the advantages of using more than one Control site, they do not agree on the ways in which this added information should be incorporated into the analyses. We expect that debate on these issues is far from over, and hope that this book serves to further the discussion and facilitate advancements in the design and application of BACI-type studies.

The Organization of This Book

The issues highlighted above are tackled directly in the second section of this book, which is devoted to elaboration of the application and design of BACI-type studies. However, prior to the implementation of any Field Assessment a number of initial issues must be considered, and some of these are highlighted in the book's first section. For example, the general goal and purpose of the study is critical, and Stewart-Oaten's first chapter (Chapter 2) tackles the standard "*P*-value culture" that places undo emphasis on the detection of impacts, rather than estimation of their magnitude or importance. Ecological parameters must also be selected for study, and although the use of bio-indicators has been often criticized, Jones and Kaly (Chapter 3) point out that any study necessarily must select a limited number of parameters from the myriad available (thus necessitating the selection of a subset of "bio-indicators"). Once appropriate species (or parameters) are selected, sampling error can constrain our ability to discern the temporal dynamics of populations, and thus impair our ability to use time series analyses to assess ecological change (Thrush, Hewitt, and Pridmore, Chapter 4). Variability also limits the power of statistical tests of impacts, and Mapstone (Chapter 5) suggests a novel way to incorporate such a constraint directly into the permitting process by simultaneously weighting Type I and Type II errors in assessment studies.

The second section of the book (*Improving Field Assessments of Local Impacts: Before-After-Control-Impact Designs*), provides the core of the book and elaborates on the theory and application of BACI-type designs. Osenberg, Schmitt, Holbrook, Abu-Saba and Flegal (Chapter 6) provide a segue from Mapstone's discussion of statistical power by evaluating sources of error in BACIPS designs and specifically evaluating the power to detect impacts on chemical-physical vs. biological (individual-based vs. population-based) parameters. Stewart-Oaten (Chapter 7) follows with a theoretical treatment of BACI-type designs, which extends and generalizes much of the earlier research on BACI(PS). Bence, Stewart-Oaten and Schroeter (Chapter 8) apply this more general and flexible BACIPS design to data derived from an intensive study of the impacts of a nuclear power plant. In the final chapter in this section (Chapter 9), Underwood offers an alternative approach in which multiple Control sites are used to detect impacts in a different way than proposed by the other authors, and which potentially can detect a greater variety of impacts (e.g., pulses as well as sustained impacts).

While BACI-type designs offer great potential to detect impacts of local perturbations, they require sampling of a Control site(s) (a site(s) sufficiently close to the Impact site(s) to be influenced by similar environmental fluctuations, but sufficiently distant to be relatively unaffected by the disturbance). In many cases, such a control does not exist, or the significant biological effects of interest are dispersed over large spatial scales and therefore are difficult (if not impossible) to detect. In such cases, a BACI-type design or other empirical measure of impact is unlikely to be able to quantify the effect with the desired level of precision. Therefore, we must be able to extrapolate results obtained from localized or smaller scale effects to those arising on larger spatial scales. This is the theme for the book's third section. Raimondi and Reed (Chapter 10) discuss how the spatial scales of impacts on chemical–physical parameters might differ from the scale of impacts on ecological parameters based on life-history features of the affected organisms. Larval dispersal is central to many of their points. Understanding the coupling between larval pools and benthic populations and their response to impacts will require integration of oceanographic models and ecological studies (Keough and Black, Chapter 11). Oceanographic processes may also "collect" larvae and pollutants in particular sites (e.g., along linear oceanographic features), and this aggregation of larvae in high concentrations of pollutants might amplify deleterious effects of many types of discharge (Kingsford and Gray, Chapter 12). Ultimately however, effects on larvae need to be translated into consequences at the population level, and in Chapter 13, Nisbet, Murdoch, and Stewart-Oaten provide an approach intended to provide an estimate of how local larval mortality, induced by a nuclear power plant, may impact the abundance of adults assessed at a regional level. Their work points out the critical need to better understand the role of compensation in the dynamics of fishes and other marine organisms affected by anthropogenic activities.

In the final section of the book, we return to the issue of Predictive vs. Postdictive approaches, emphasizing how and why the Predictive phase (which typically yields an Environmental Impact Report or Statement) should be integrated with the Postdictive phase (which yields the Field Assessment) to improve the overall quality of the entire process. The introductory chapter in this section (Chapter 14: Schmitt, Osenberg, Douros, and Chesson) provides a brief summary of the current state of the EIR/S process in the United States and Australia, and Carney (Chapter 15) evaluates the biological data that have been collected with regard to EIR/S studies (as well as Field Assessments). He concludes that numerous problems (including taxonomic errors, design flaws, statistical inaccuracies) plague even the most extensive studies. Piltz (Chapter 16) then elaborates on how institutional constraints can impede sound scientific investigations that require long-term monitoring. His chapter encourages both scientists and administrators to find solutions that ensure the research continuity that is required to obtain the most defensible Field Assessments. If EIR/S consistently fail to yield consensus, or if too few data exist to determine the accuracy of such studies, considerable debate can ensue. This often leads to judicial involvement in the EIR/S process, which in turn leads to tremendous effort expended on documentation during the EIR/S process, but with little additional clarity regarding the likely or actual environmental impacts of a project (Lester, Chapter 17).

Ultimately, the interplay between the Predictive process (EIR/S or Risk Assessment) and the Postdictive process (the Field Assessment) is critical to help guide our development of frameworks used to predict and understand anthropogenic impacts: are our predictions accurate, and if not how might we modify our approach? Ambrose, Schmitt and Osenberg (Chapter 18), provide an audit of an intensive study of the San Onofre Nuclear Generating Station (SONGS) and compare effects that were predicted during the EIR/S process with those subsequently observed. Their findings reveal that the EIR/S process (even in such a major project) revealed little of the actual impacts; even a detailed scientific study conducted by an independent committee erred in a number of crucial ways. Their results demonstrate the need for continued vigilance in conducting well-designed monitoring studies, such as those using BACI-type analyses.

At times the tone of many of these contributions is rather critical of existing approaches and even the new approaches outlined in other chapters. It is only through healthy debate of the merits of alternative designs, and by better integration of administrative and scientific goals, that improvements in the EIR/S and Field Assessment processes occur. Indeed, despite this body of criticism, it is undeniable that refinements in our scientific tools have led to recent improvements in our understanding of anthropogenic impacts. This can only continue by avoiding complacency and by continuing to develop and refine new tools that can be used to tackle these important issues. Of course, realization of our ultimate goal depends upon expanding our basic knowledge of the dynamics and functioning of ecological systems, understanding the mechanisms by which anthropogenic activities impact these systems, and incorporating this information into models and theory to permit us to predict the occurrences of future impacts. But first, we must be better able to quantify the actual impacts that specific activities have induced in ecological systems. This simple goal is neither trivial nor commonly realized, but it is fundamental and achievable. Our hope is that this book helps to further advance understanding of the interactions between human activities and our impacts on our environment.

Acknowledgments

We gratefully acknowledge the helpful comments of S. Holbrook, C. St. Mary, and A. Stewart-Oaten. Research that led to the production of this chapter was funded by the Minerals Management Service, U.S. Department of Interior under MMS Agreement No. 14-35-001-3071, and by the UC Coastal Toxicology Program. The views and conclusions in this chapter are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the U.S. Government.

References

Bartell, S. M., R. H. Gardner, and R. V. O'Neill. 1992. Ecological risk estimation. Lewis Publishers, Chelsea, Michigan.

Box, G. E. P., and G. C. Tiao. 1975. Intervention analysis with applications to economic and environmental problems. Journal of the American Statistical Association **70**:70–79.

Buckley, R. 1991a. Auditing the precision and accuracy of environmental impact predictions in Australia. Environmental Monitoring and Assessment **18**:1–24.

Buckley, R. 1991b. How accurate are environmental impact predictions? Ambio **20**:161–162.

Campbell, D. T., and J. C. Stanley. 1966. Experimental and quasi-experimental designs for research. Rand McNally, Chicago, Illinois.

Carpenter, S. R. 1990. Large-scale perturbations: opportunites for innovation. Ecology **71**:453–463.

Culhane, P. J. 1987. The precision and accuracy of U.S. Environmental Impact Statements. Environmental Monitoring and Assessment **8**:217–238.

DeMartini, E. 1987. Final report to the Marine Review Committee. The effects of operations of the San Onofre Nuclear Generating Station on fish. Marine Science Institute, University of California, Santa Barbara.

Eberhardt, L. L. 1976. Quantitative ecology and impact assessment. Journal of Environmental Management **4**:27–70.

Erickson, P. A. 1994. A practical guide to environmental impact assessment. Academic Press, San Diego, California.

Gilpin, A. 1994. Environmental impact assessment (EIA). Cambridge University Press, Cambridge, England.

Green, R. H. 1979. Sampling design and statistical methods for environmental biologists. Wiley and Sons, New York, New York.

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

Jassby, A. D., and T. M. Powell. 1990. Detecting changes in ecological time series. Ecology **71**:2044–2052.

Kastendiek, J. and K. R. Parker. 1988. Interim technical report to the California Coastal Commission. 3. Midwater and benthic fish. Marine Review Committee, Inc.

NRC (National Research Council). 1990. Managing troubled waters: the role of marine environmental monitoring. National Academy Press, Washington, DC.

NRC (National Research Council). 1992. Assessment of the U.S. outer continental shelf environmental studies program. II. ecology. National Academy Press, Washington DC.

Osenberg, C. W., S. J. Holbrook, and R. J. Schmitt. 1992. Implications for the design of environmental impact studies. Pages 75–89 in P. M. Griffman and S. E. Yoder, editors. Perspectives on the marine environment of southern California. USC Sea Grant Program, Los Angeles, California.

Petts, J., and G. Eduljee. 1994. Envionmental impact assessment for waste treatment and disposal facilities. John Wiley and Sons, Chichester, England.

Skalski, J. R., and D. H. McKenzie. 1982. A design for aquatic monitoring systems. Journal of Environmental Management **14**:237–251.

Spellerberg, I. F. 1991. Monitoring ecological change. Cambridge University Press, Cambridge, England.

Stewart-Oaten, A., W. W. Murdoch, and K. R. Parker. 1986. Environmental impact assessment: "psuedoreplication" in time? Ecology **67**:929–940.

Stewart-Oaten, A., J. R. Bence, and C. W. Osenberg. 1992. Assessing effects of unreplicated perturbations: no simple solutions. Ecology **73**:1396–1404.

Suter, G. W., II, editor. 1993. Ecological risk assessment. Lewis Publishers, Boca Raton, Florida.

Tomlinson, P., and S. F. Atkinson. 1987. Environmental audits: A literature review. Environmental Monitoring and Assessment **8**:239–261.

Underwood, A. J. 1991. Beyond BACI: experimental designs for detecting human environmental impacts on temporal variations in natural populations. Australian Journal of Marine and Freshwater Research **42**:569–587.

Wathern, P. 1988. Environmental impact assessment. Unwin Hyman Ltd., London, England.

Westman, W. E. 1985. Ecology, impact assessment, and environmental planning. John Wiley and Sons, New York, New York.