The Efficiency of Direct Public Involvement in Environmental Policymaking: An Experimental Test

Bruce, C. [1] and J. Clark[2]

[1] Department of Economics, University of Calgary, Canada
[2] Department of Economics, University of Canterbury, New Zealand
Email: jeremy.clark@canterbury.ac.nz

Keywords: *Public policymaking, Nash Bargain, experiments*

EXTENDED ABSTRACT

In one of the most ambitious forms of environmental decision-making, representatives of interested parties – environmentalists, developers, farmers, loggers, miners, etc. - are charged with the responsibility of developing a set of public policies that is acceptable to all of them. Although this approach has become increasingly popular, and has been widely discussed in the academic literature, little is known about the characteristics of the outcomes that are reached in this type of negotiation. We do not know, for example, whether these outcomes meet the standard criteria for efficiency or equity.

In this paper, we use laboratory experiments to test whether a number of *axiomatic* models of bargaining can predict the behavior of the parties to environmental decision making. In recognition of the multi-dimensional aspect of most public land use conflicts, we ask pairs of subjects to negotiate over two goods, without the possibility of cash side payments. We thus provide one of the first experimental tests of a prediction associated with the Edgeworth Box: that parties with an initial endowment that is Pareto inefficient will make trades until they reach a Pareto efficient allocation (See Figure 1). We further test whether parties in particular reach the Nash bargain when it coincides with, is orthogonal to, or conflicts with outcomes that maximise the parties' joint payoffs. In the latter case, we also examine the effect of providing parties with full or partial information regarding payoffs. In all cases, the Nash bargain is set to also equalize the parties' payoffs.

We find that bargaining pairs were drawn to Pareto efficient allocations, and the Nash bargain in particular under all treatments with full information as to payoffs. This was strongest when the Nash bargain coiicided with the joint-payoff maximum, but also evident when it was orthogonal to the joint maximum, or even when it diverged. When the pairs were provided only with information about their own payoff functions, however, agreements only "approached" the Nash bargain over progressive rounds, with slight deviations that preserved some of the initial advantage of one of the parties.

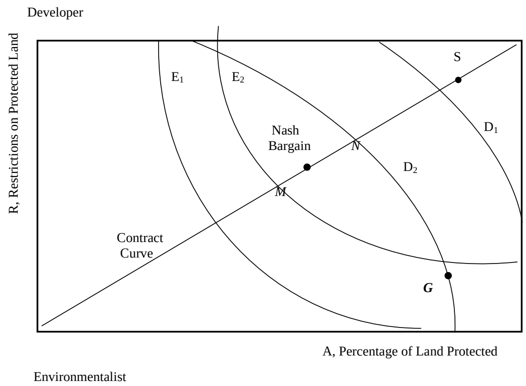

Figure 1: An Edgeworth Box representation of land use conflict

1 INTRODUCTION

As citizens have become better informed and educated, they have sought increased levels of accountability from their governments. This has been particularly true in the field of environmental policy-making. In the last three decades, citizens have demanded that they be given ever-greater involvement in the setting of policies concerning air and water pollution, endangered species, and public land use. Whereas governments were initially able to meet these demands by soliciting views through public hearings, they have felt increasing pressure to allow citizens to become directly involved in the processes by which rules, regulations, and government policies are drafted.

Hundreds of different types of direct involvement have sprung up at virtually every level of government in North America and Europe. These processes have appeared under a number of different guises – among them consensus-building, mediation, conflict resolution, environmental dispute settlement, negotiated rulemaking, and collaborative management – reflecting their varied origins. But all share the properties: that they are open to all interested parties, that those parties are expected to develop policies that are acceptable to all participants, and that the parties operate under the assumption that their proposals will be adopted by the government.

A large literature has investigated examples of public policymaking in case study form; and the advantages and disadvantages of the various forms of that involvement have been extensively debated. What this literature lacks, however, is a framework within which predictions can be made concerning the characteristics of the negotiated outcomes. Can public policymaking outcomes be expected to be efficient, for example? Will they be fair or equitable? Will they maximise social welfare, regardless of how the backstop policy is set? The purpose of this paper will be, first, to develop such a framework, based on the classic Edgeworth Box model; and, second, to use laboratory experiments to test some of the predictions associated with that model.

The paper is divided into five sections. In Section 2, we argue that any model of the manner in which environmental policies are negotiated must be multidimensional. We develop a model to pose several predictions concerning the outcomes that negotiators can be expected to reach. Section 3 describes our experiment, which we use to test our predictions. In Section 4, we report our results. Section 5 concludes with some implications of our findings.

2. A MODEL OF NEGOTIATION

In most cases, *public* environmental policymaking involves the development of plans to manage publicly owned natural resources. Examples include decisions about public lands, such as national forests (Terhune, 1998; Tableman, 1999); wetlands (State of Minnesota, 1997) or watersheds (Knudson, 1999). Universally, these decisions are multi-dimensional.

The multi-dimensionality of such problems argues for the use of a bargaining model in which negotiators are required to make a number of decisions simultaneously, and thus “trade” with one another. For this purpose, we develop an Edgeworth Box model involving the government, two parties, and two characteristics. We then employ that model to illustrate several predictions from axiomatic bargaining theory that might be expected in public policy making.

For expository purposes, suppose that the government wishes to allocate a parcel of public land between commercial uses and public uses. Within those allocations, various restrictions may be placed on permissible uses: for example, ranchers might be required to prevent their cattle from disturbing riparian ecosystems. We assume that these alternatives can be captured by two characteristics: (i) the number of acres of public land to be set aside as environmental reserve, A, and (ii) the level of (environmental) restrictions to be placed on the commercial and recreational use of each acre of land placed in that reserve, R.

There are three actors in our model: the *government* and two interest groups, *environmentalists* and *developers*. The government’s goal is assumed to be the maximization of the sum of the interest groups’ utilities. Its ability to achieve this, however, is constrained by its lack of information about the parties’ true utility functions. Accordingly, the government uses its best estimate of the parties’ preferences to select a backstop policy G, $(= G(A_g,R_g))$, but recognizes that G may be Pareto inefficient. In an effort to identify a superior outcome, it establishes a process in which environmentalists and developers are invited to construct their own proposal that the government would then implement. It is assumed that G will be adopted if the parties fail to reach a joint agreement. Finally, we assume that the parties are constrained to bargain only over A and R: there is no third commodity such as money that either party could use to make “side payments.” For example, we assume that logging companies are unable to pay environmental groups to accept policies that are environmentally harmful.

Environmentalists and developers thus bargain with one another "in the shadow" of policy *G*. Figure 1 above, a conventional Edgeworth Box, represents the preference functions of the two parties to this bargaining process. As environmentalists prefer both additional acres of reserve and stronger restrictions on allowable uses of those acres, their indifference curves are convex to the origin and increasing in utility with both *A* and *R*. Conversely, the developers' indifference curves are concave to the origin and increase in utility as both *A* and *R* decrease. The set of outcomes that is Pareto superior to *G* is represented by the bargaining lens between the indifference curves of the two parties that intersect at *G* in Figure 1. The set of outcomes that is Pareto efficient is given by the portion of the contract curve within the bargaining lens, or *MN*.

Assuming that the government's objective is to maximize the sum of the parties' utilities the social optimum must also be Pareto efficient and will lie on the contract curve. However, that optimum need not fall within the bargaining lens. If G has been chosen inappropriately, it is possible that one of the parties may be worse off at the optimal allocation than at the backstop allocation. One such outcome has been identified as S in Figure 1.

Ideally, the government would like to know whether the parties will be able to bargain to an agreement and, if so, whether it will (i) represent an improvement on *G* and (ii) approach the social optimum, S. Four predictions concerning the negotiated outcome are commonly made in the bargaining literature:

Pareto superiority: In standard versions of the Edgeworth box model, it is predicted that the bargaining parties will reach a Pareto superior outcome to the backstop position. That is, the parties will bargain to an outcome inside the bargaining lens formed around G. Note that, in the absence of side payments, this excludes the possibility that the parties will reach a joint utility-maximising outcome outside the bargaining lens, such as S in Figure 1.

Pareto efficiency: More specifically, it is usually argued that, in the absence of impediments, the negotiated outcome will be Pareto efficient. In Figure 1, if G is the backstop allocation, the parties will agree to an outcome on *MN*.

The Nash bargain: John Nash (1950) studied bargaining games in which each party was assumed to be both individually rational and well-informed about the opponent's utility function. In these cases, he showed that if the parties' behavior satisfied four intuitive axioms, they would choose an outcome that maximised the product of their respective net gains relative to the endowment point *G*.

Equality of outcomes: Experimental studies often find that at least some individuals seek allocations that divide benefits in such a way as to reduce or eliminate inequality of final outcomes.

The experiment that we report in Sections 4 and 5 is designed to test these predictions. We test whether our subjects are able to negotiate Pareto efficient agreements, particularly the Nash bargain, when the latter is identical to, orthogonal to, or divergent from the joint payoff maximising outcome. In the last case, we also test the effect of moving from complete to private information about the payoffs the parties receive at each allocation.

3. EXPERIMENTAL DESIGN

3.1 Design Features Across All Treatments

We designed an experiment in which subjects were recruited in groups of 10, and each given an induced value payoff function over two dimensions of abstract goods, *X* and *Y*. Half of the subjects were assigned one payoff function, and the other half a second payoff function. We shall refer to the two preference types induced by these payoff functions as environmentalists, *E*, and developers, *D*, though neutral labels were used in the experiment. We selected Cobb Douglas payoff functions for the two types to generate convex indifference curves:

$$P_E = A_E X_E{}^a Y_E{}^a + B_E \qquad (1)$$

$$P_D = A_D X_D{}^a Y_D{}^a + B_D \qquad (2)$$

Each type of individual, *i*, was endowed with an initial backstop allocation of $X_{i,G}$ and $Y_{i,G}$. In order to offer our subjects sufficient numbers of choices that they would not be implicitly "directed" to the outcomes we predicted, we set the sum of $X_{D,G} + X_{E,G}$ and the sum of $Y_{D,G} + Y_{E,G}$ each equal to twenty. We next chose a non-symmetric endowment backstop of $(X_{E,G}, Y_{E,G}) = (18, 5)$ and $(X_{D,G}, Y_{D,G}) = (2, 15)$. In all treatments, this implied that the contract curve was located roughly between $(X_E, Y_E) = (9.5, 9.5)$ and $(X_E, Y_E) = (14.5, 14.5)$.

Because risk preference is thought to influence bargaining outcomes (Murnighan, Roth and Schmeidler, 1988), subjects' risk attitudes were elicited prior to the bargaining instructions using the method of Holt and Laury (2002).

After studying their own payoff tables (and those of their opponents in the full information treatments) for some time, subjects were then placed together in pairs, one environmentalist with one developer, and allowed a two minute period of unstructured communication in which they might agree to any alternative allocation of *X* and *Y*. To be accepted as valid, negotiated outcomes had to be technically feasible, or

$$X_E + X_D \leq X_{E,G} + X_{D,G} = 20 \qquad (3)$$

$$Y_E + Y_D \leq Y_{E,G} + Y_{D,G} = 20 \qquad (4)$$

To successfully register a negotiated outcome other than the backstop, one of the bargaining pair had to describe the allocation on a form, and the other had to tick a box signifying agreement.

To control for the effects of accumulating income on risk preference, only one of the five rounds was implemented at the end of experiment, chosen by the throw of a die. To prevent subjects from being able to make credible offers of side payments after the experiment, a different random draw was carried out for each person in privacy when being paid to determine which round to count.

Our mixing protocol ensured that each member of one type was paired serially with all five members of the other type during the five rounds. The experiment was conducted manually.

Each pair's conversation was recorded with a micro-cassette player located midway between them to one side of the tables. We hope that the "uncontrolled aspects of social interaction" (Roth 1995) introduced by unstructured face-to-face communication were more than compensated by the increased parallelism between our design and actual face-to- face negotiations that take place between stakeholders in public policy making.

3.2 Design Features of Each Treatment

In all treatments, we chose the *A*'s, *B*'s and *a*'s of the payoff functions in such a way as to keep constant the following:

1. box size $X_E + X_D = 20$ and $Y_E + Y_D = 20$
2. endowments $(X_{EG},Y_{EG})=(18,5),(X_{DG},Y_{DG})=(2,15)$
3. Nash bargain $(X_E,Y_E) = (12,12)$, $(X_D,Y_D) = (8,8)$
4. sum of endowment values:

$$A_E 18^a 5^a + B_E + A_D 2^a 15^a + B_D = \$20.00$$

5. sum of Nash bargain values:

$$A_E 12^a 12^a + B_E + A_D 8^a 8^a + B_D = \$36.20$$

In addition, we set the parameters to ensure that the value of the joint payoff (i.e. the sum of the parties' payoffs) was substantially higher at the NB than at G and, where divergent, was higher than the NB. Finally, the parameters were adjusted in such a way that, across five experimental treatments, increasingly stringent tests were placed on the hypothesis that the parties would select the Nash bargain.

In Treatment I, we chose the parameters in such a way that (i) the joint payoff was maximised at a single outcome, (developers receive (12, 12)/environmentalists receive (8, 8)), and (ii) the outcome at which the joint payoff was maximised was also the Nash bargain *and* a point at which the money values of the payoffs to the parties were equalized (at \$18.10 each). To simplify the presentation, subjects were provided colored payoff tables showing the specific earnings that both they and their opponent would receive from the experimenter for all feasible combinations of *X* and *Y*. Note that, as we did not include negative values in the payoff tables, some extreme allocation cells in each table did not contain entries.

In one sense, it might not have been difficult for the parties to reach the Nash bargain in Treatment I because they only needed to identify the (unique) outcome at which their joint payoff was maximised. Even here, however, the challenge facing the parties was much more complex than that presented to subjects in the variable sum experiments that were described in Section 3, as subjects now had over two hundred outcomes from which to choose, instead of the eight or so that were offered in previous experiments.

In Treatment II we set the parameters in such a way that the joint payoff at every point on the contract curve now equaled the maximum – in order to address the possibility that the unique joint payoff maximum in Treatment I may have acted as a focal point. As in Treatment I, however, the parties received equal payoffs at the Nash bargain and both parties were given full information about their opponent's payoffs.

In Treatment III, we wished to investigate the effect of separating the joint payoff maximum from the Nash bargain. Our goal was to determine whether the parties might be drawn towards the "social optimum" (joint payoff maximising allocation) when the bargaining lens associated with the endowment did not include this allocation. The parameters chosen for this treatment left the individual payoffs equal to one another at the Nash bargain (as in the first two treatments) but moved the maximum joint payoff to a point in the "northeast" portion of the contract curve. Again, the parties were given full information about one another's payoff tables.

Recognizing that subjects might also care about initial inequality, ideally we would have liked them to have started with identical endowment incomes across all three treatments. This turned out to be infeasible in Treatment III, and we settled for the second best of setting the *sum* of endowment earnings equal to the same value as had been used in the first two treatments, $20. This resulted in an endowment income advantage for the developer ($13.15 vs. $ 6.85), unlike in Treatments I and II where endowments were roughly equal.

To test for the effect of information about the opponent's preferences, Treatment IV reproduced Treatment III in all respects, except that subjects were given only their own type's payoff table. They were instructed verbally that they were under no obligation to share the information on their payoff table with the other party during negotiations, but were given the option to reveal or not reveal that information as they chose.

To control for preferences for final equality, the Nash bargain occurred at an outcome at which the parties received equal payoffs ($18.10 each) for all treatments. This will likely have the effect of consistently raising support for the Nash bargain above what it would otherwise be, but preserve our ability to make comparisons across treatments. Finally, subjects in all five treatments completed a questionnaire supplying demographic information and explaining what they tried to achieve during bargaining.

4. THE RESULTS

Seventeen sessions with 10 subjects each were run at the University of Canterbury. Four sessions were run per treatment, with an additional session run for Treatment IV that will be described shortly. Thus each treatment contained 40 people who provided 20 paired bargaining outcomes per round over five bargaining rounds. Each bargaining outcome consisted of the physical allocation of X and Y between an Environmentalist and Developer, and their resulting respective earnings. Each session took roughly 90 minutes, and subjects earned on average approximately 25.00NZ$ (1.00NZ$ = 0.72US$).

A complication arose with the intended final (fourth) session of Treatment IV. A subject who had participated previously in a session of Treatment I (where the Nash bargain is a prominent solution) participated again in this session, where the Nash solution was not as prominent. This subject, and every person he was paired with for a given round, agreed to settle at the Nash bargain. No other bargaining pair in this session chose to settle at that point on any round, including those who had agreed to it when paired with this individual on previous rounds. The potential contamination of this session's results meant that a replacement final session of Treatment IV was run. Its results were similar to those of the pairs that excluded the repeat participant in the contaminated session. We thus report the results for Treatment IV using the four "clean" sessions. But given the robustness of bargaining outcomes in the contaminated session that did not directly involve the repeat subject, we also report augmented results for Treatment IV that include that session's four pairs of outcomes for each round that exclude the repeat subject.

We divide our discussion of the results from these experiments into five sections. In the first of these we discuss whether the parties reached agreement. In the second and third, we ask whether agreements were, respectively, Pareto superior to the backstop and Pareto efficient. In the final two we ask whether the parties chose the Nash bargain, and whether they maximised the sum of the parties' payoffs.

4.1 Agreement Rates

When averaged over all five rounds, agreement rates were roughly similar across the first four treatments: 92% in Treatment I, 88% in Treatment II, 90% in Treatment III, and 82% in Treatment IV. Nevertheless, panel random effects logit analysis of whether the parties reached agreement confirms that agreement rates in Treatment IV were significantly lower than in the other treatments. Upon closer inspection, this lower agreement rate was due almost entirely to a sharply lower agreement rate in Treatment IV's Round 1

(40%), as subjects with private payoff information struggled to reach agreements within the first two minute round. Excluding the first round, the overall agreement rates for Treatments I, II, III and IV were much more similar, at 97.5%, 92.5%, 95.0% and 92.5%, respectively.

4.2 Pareto Superiority

The first prediction of the axiomatic models (Section 2) is that the bargained outcome will be Pareto superior to the backstop position; that is, that subjects will choose an outcome within the bargaining lens. This prediction was amply supported across all four treatments. Counting pairs who did not reach agreement as within the lens, the overall rates across all five rounds for Treatments I, II, III and IV were 100%, 97%, 98% and 99%, respectively. For Treatments III and IV these results do more than confirm elementary rationality: as, in those treatments, the parties could have gained as much as $11.17 (over the Nash bargain) by settling outside the bargaining lens. This confirms our hypothesis that, in the absence of side payments, if the government sets a poor backstop policy that creates a bargaining lens excluding the joint payoff maximum, parties will not go there.

4.3 Pareto Efficiency

The prediction of Pareto efficiency - that the parties will bargain to the contract curve within the bargaining lens - received strongest support in Treatment I, strong support in II and III, and at least some support in IV. Treating lack of agreement as a Pareto inefficient outcome, the rates of Pareto efficiency averaged across all five rounds for Treatments I, II, III and IV were 86.0%, 71.0%, 73.0% and 27.0%, respectively. Excluding the first bargaining round, when disagreement rates were highest, agreement rates were 97.5%, 77.5%, 78.8% and 31.3%, respectively.

In formal statistical tests, 95% exact binomial confidence intervals around the sample proportion of pair outcomes on the contract curve contained the proportion 100% in three of the four final rounds for Treatment I, but not in any round for the other treatments. So we can reject the hypothesis that *all* subjects reached the contract curve within the lens for any but Treatment I. To compare results across treatments, we use random effects logit panels for "reached/did not reach the contract curve". We find subjects were not significantly less likely to reach the contract curve in Treatments II or III than they were in Treatment I, but were less likely to do so in Treatment IV, where information was private.

Nevertheless, there is strong evidence from subject earnings that agreements closely approached the contract curve as the subjects gained experience over the five rounds of bargaining. Mean payoffs tended to converge on the Nash bargain/equal-payoff outcome (at which both parties obtained $18.10). By the fifth round total payouts did not deviate significantly from that outcome in any treatment. In the final round of Treatment IV, for example, nineteen of twenty-four allocations were on, or within one unit of, the contract curve - nine having settled at the Nash bargain, (12,12)/(8,8), and five having settled at each of (12,11)/(8,9) and (11,12)/(9,8).

This is a remarkable finding. When faced with two payoff tables, each of which contained over two hundred cells, the parties were able to reach one of a handful of efficient (or approximately efficient) outcomes – even when those outcomes did not maximise their joint payoffs and even when the parties were not provided with information about their opponents' payoff functions. These results suggest strongly that even in the face of complex issues, parties are able to identify and negotiate efficient agreements.

4.4 Nash Bargain

The specific prediction that parties would reach the Nash bargain also receives support, although not as strong as that received by Pareto efficiency. Averaged across all five rounds, the Nash bargain was selected in 81.0%, 70.0%, 69.0% and 22.0% of Treatments I, II, III and IV, respectively. Averaged across the last four rounds, the rates were 92.5%, 77.5%, 75.0%, and 27.5%. In formal statistical tests, 95% exact binomial confidence intervals around the sample proportion of pair outcomes at the Nash bargain contained the proportion 100% only in rounds three and five of Treatment I, and never for the other treatments. So we can reject the hypothesis that *all* subjects reached the Nash bargain, particularly for Treatments II-IV. At the same time, however, by the final round the negotiated outcomes did not result in earnings that were statistically different from the Nash bargain, at the five percent level in any of the treatments.

To comparing outcomes *across* treatments, we ran random effects logit panels for reaching the physical Nash bargain. Subjects were not significantly less likely to reach the Nash bargain in Treatments II and III than they were in Treatment I, but were less likely to do so in Treatment IV, where information was private. Running random effects panels for deviations of actual earnings from Nash levels, we find that only in Treatment IV was the differential in total earnings between the negotiated outcome and the Nash bargain significantly greater than it had been in Treatment I. Despite the lower overall earnings in Treatment IV, by the later rounds the most common outcome was the Nash bargain or one of the two allocations immediately adjacent to it. It should be noted that the allocations immediately adjacent to the Nash preserved a slight advantage for the developer, who had started the bargaining round with the $13.15/$6.85 advantage.

4.5 Maximization of Joint Payoffs

We designed Treatments III and IV to test whether support for the Nash bargain would diminish when it did not coincide with the joint payoff maximum. In each of these treatments, the sum of the parties' payoffs would be maximized if the Developer were to agree to an outcome at which it obtained fewer units of both goods than it was able to obtain at the backstop. The sum of the parties' payoffs at the joint maximum (our measure of the social optimum), $47.37, significantly exceeded the sum available at the Nash bargain, $36.20.

We ran random effects panel regressions of deviations in earnings from the Nash Bargain across treatments. With full information maintained, the earnings outcomes in Treatment III were unaffected by the presence of a joint maximum that differed appreciably from the Nash bargain. Indeed, the outcomes reached in that treatment lay just as close to the Nash bargain as in Treatments I and II.

Second, as already mentioned, in the private information treatment the negotiated outcomes were more likely to deviate from the Nash bargain than they had been in the three public-information treatments. Importantly, however, these deviations did *not* tend "towards" the joint maximum. Rather, average total payoffs in Treatment IV rose to $35.88 (within one percent of the total available from the Nash bargain) by the fifth round of bargaining, and only one of the 120 bargaining pairs agreed to the joint maximum. That is, deviations from the Nash bargain in Treatment IV appear to have been related to the parties' lack of information, or willingness to preserve relative income rankings, and not to the differences between the Nash bargain and the joint maximum.

5. SUMMARY AND CONCLUSIONS

We have set out to represent the process of public policymaking over multiple dimensions of public land use as a cooperative bargaining problem. We have designed and run laboratory experiments to test whether the predictions of standard axiomatic bargaining models – Pareto superiority, Pareto efficiency, and Nash bargaining – might successfully predict the outcomes of such public policymaking negotiations.

We sought to create an analogy to the case in which a government sets a backstop policy prior to negotiations that the stakeholders must live with if negotiations fail. Under full information we tested whether parties would reach the Nash bargain when it coincided with, was orthogonal to, or diverged from the socially optimal allocation. We repeated the last, most difficult treatment with the more realistic case of private information.

Even with the small stakes, limited time and absence of side payments, we find encouraging evidence that subjects reach Pareto superior and efficient outcomes, and reach or converge to the Nash bargain under all treatments. It remains for future work to see whether the Nash bargain will perform as well when it does not also equalize final earnings.

7. REFERENCES

Holt, C., and S. Laury (2002) "Risk aversion and incentive effects," *American Economic Review*, 92 (5), 1644-55.

Knudson, T.(1999) "The Ranch Restored," *High Country News*, 31(4), (March 1), 13-16.

Rubinstein, A. (1982) "Perfect Equilibrium in a Bargaining Model," *Econometrica*, 50 (1) 97-110.

State of Minnesota (1997) *Minnesota Wetlands Conservation Plan, Version 1.0*, (Department of Natural Resources, St. Paul, Minnesota)

Stout, R., D. Decker, B. Knuth, J. Proud, and D. Nelson (1996) "Comparison of Three Public-Involvement Approaches for Stake-holder Input" *Wildlife Society Bulletin* 24(2), 312-317.

Tableman, M. (1990) "Case Study 1: San Juan National Forest Mediation" in Crowfoot, J., and J. Wondolleck, *Environmental Disputes*, (Washington, D.C., Island Press), 32-65.

Terhune, P., and G. Terhune (1998) *QLG Case Study*, (Quincy Library Group), www.qlg.org/pub/miscdoc/terhunecasestudy.htm (accessed August 18, 2003)