The Canadian Journal of Program Evaluation Vol. 10 No. 2 Pages 113–122
ISSN 0834-1516

SENSITIVITY ANALYSIS IN OUTCOME EVALUATIONS: A RESEARCH AND PRACTICE NOTE

Ken Watson
Rideau Strategy Research Associates
Ottawa, Ontario

Abstract: Every evaluation study uses data that are uncertain to some degree. Therefore, the analyst and the decision-maker need to know how much the outcome of the evaluation varies given the plausible variation in uncertain data inputs. That is, how sensitive is the outcome of the analysis to a particular input variable? This note discusses what characteristics make for sensitivity, what techniques to use to clarify sensitivity (especially graphic techniques), and how to interpret the result.

Résumé: Chaque étude utilise de l'information qui est incertain à un certain degré. Alors, l'analyste et celui qui prend les décisions ont besoin de connaître combien le résultat de l'évaluation est inconstant. Cela veut dire, dans quelle mesure est le résultat de l'analyse sensible à une entrée variante particulière. Cette note discute quelles caractéristiques comporte la sensibilité, quelles techniques à utiliser pour identifier la sensibilité (particulièrement les techniques graphiques), et comment interpréter le résultat.

Some evaluation studies, although not all, measure the outcomes of a program by a quantitative calculation (a "model"). An example is the "cost-benefit" model, in which discounted net costs are balanced against discounted net benefits to obtain the "net present value," or NPV. In such models, the outcome is typically influenced by several uncertain factors. This is true in fields as diverse as health, education, employment, or economic development (Baird, 1989). When the evaluator is dealing with such a model, an analysis of how "sensitive" the outcome is to changes in those uncertain factors is a useful first step toward two things: identifying where it is worth spending money to obtain more precise data, or where practical action can be taken to limit uncertainty (for example, by redesigning program components or simply keeping a watchful eye when managing the program); and communicating to decision-makers the extent of the uncertainty and risk in the program.

This note sets out some of the techniques (mainly graphic) that are part of the sensitivity analysis "tool kit." The example considered is sensitivity analysis in a cost-benefit decision by a hospital whether to buy a heart-lung machine.

Sensitivity analysis is a limited tool. It treats variables one at a time, holding all else constant. The simultaneous actions and interactions among variables in the real world are not dealt with. A variable that appears to be key when considered in isolation might or might not be key when considered along with other variables that strengthen or weaken its effect on the outcome of the program. Only a risk analysis (Hertz & Thomas, 1984) can accurately identify the true influence of the variable. Nevertheless, sensitivity analysis is helpful in exploring the "deterministic" model, which tends to be the second phase of a general analysis:

1. Build a deterministic model using single "best" values (base values) for the input variables.
2. Explore the outcome's sensitivity to each input variable, and take action.
3. Make a full risk analysis using probabilities for many variables simultaneously.

Sensitivity analysis is important for the analyst starting to understand the model. As this understanding develops, the analyst can take action when appropriate. In some cases the only action that can be taken is to obtain better data. For example, the outcome of a decision whether to purchase a heart-lung machine is sensitive to "probability of an influenza epidemic," which is not controllable by the analysts. In other cases, practical action might be taken to fix or constrain the value of the variable. For example, if the outcome is particularly sensitive to an operator's wage rate, then this rate might be negotiated beforehand. Fixing the wage rate would dramatically lower the sensitivity of the outcome to this variable. The more the sensitivities can be minimized, the more precise will be the estimate of the outcome.

Gross Sensitivity

In its simplest form, which we might call "gross sensitivity," sensitivity analysis involves calculating, one at a time, how much the net present value changes if influencing variables change by a standard percentage, say 10%. Consider the purchase of a heart-lung machine whose net present value is affected by four variables: insurance costs, operating costs, the price of the machine, and the usage rate (Table 1). A quick glance indicates that the decision whether to purchase the machine is quite sensitive to three of the four variables.

What Determines Sensitivity?

A superficial interpretation of Table 1 could be misleading. One should keep in mind that the "effective sensitivity" of the outcome to a particular input variable is determined by four factors:

1. How *responsive* is the net present value to changes in the variable?
2. How large is the variable's *range of plausible values*?
3. How *volatile* is the value of the variable?
4. Can the range or the volatility of the variable be *controlled*?

The first of these factors, the responsiveness of the NPV to the input variable, has two components. The first component is the direct influence of the variable on the NPV. The second is the indirect influence of the variable, through its relationships with other variables that are themselves related to the net present value. Positive correlations with other influential variables will magnify the ultimate influence of both, and negative correlations will decrease their influence. These influences can be fully identified only when one has set up a simulation model that is capable of dealing with the simultaneous interactions of many variables.

Table 1
Gross Sensitivity of NPV to Four Input Variables

Variable	Change in NPV in Response to 10% Change in Variable
Insurance costs	15%
Operating costs	21%
Machine price	7%
Usage rate	19%

Sensitivity and Decision-making

We are most interested in the sensitivities that might change a positive decision on the project to a negative decision, and vice versa. Four calculations help us to estimate the likelihood of such a switch:

1. What is the *range of influence*? That is, how much does the NPV change when the variable changes from its lowest to its highest plausible value?
2. Does this range of influence contain a zero NPV? If it does, then the variable has a *switching value*, that is, a value at which our appraisal of the project switches from positive to negative.
3. What is the *switching ratio* for the variable? That is, by what percentage does the variable have to change to hit a switching value?
4. What is the *switching probability*? That is, how likely is the variable to reach the switching value?

If we add this information to the gross sensitivity calculation shown in Table 1, we begin to have a reasonably complete picture of the likely sensitivity to a variable, although, of course, within the limits imposed by considering one variable at a time (see Table 2).

Scanning all of the data in Table 2, one can see that the NPV is sensitive to neither insurance costs nor machine price. In particular, neither variable can move the NPV enough to hit a switching value. In contrast, NPV is sensitive to both operating costs and usage rate. Of these two, usage rate is obviously more volatile: although

Table 2
Several Indicators of Single-Variable Sensitivity

	Variables			
	Insurance costs	Operating costs	Machine price	Usage rate
Gross sensitivity	15%	21%	7%	19%
Range of influence	10%	17%	5%	35%
Switching value	no	yes	no	yes
Switching ratio	—	9%	—	63%
Switching probability	—	40%	—	42%

it has to change by a much larger percentage to hit a switching value, it is about equally likely to do so (40% compared with 42%). In other words, although usage rate has to change more than operating costs to cause a crucial change in the net present value, its volatility makes it equally likely to do so. On this evidence, we would tentatively conclude that the two key variables are equally influential.

Two-Input-Variable Sensitivity Analysis

So far we have analyzed sensitivity one variable at a time. We can extend this to two variables as the next step toward true risk analysis. Scenarios defined by two interacting variables, although still not complete and realistic indicators of sensitivity, are at least an improvement on single-variable analysis. The joint influences of two input variables can be described in a matrix like the one in Table 3 below.

From Table 3 one can see, for each cost of capital, what usage rate has to be attained if the machine is to be economical. In fact, one can see that a line that gives all of the combinations of discount rate and usage rate resulting in an NPV of zero could be drawn across the table, dividing it into two "strategy regions" (defined by combinations of discount rates and usage rates). In one of these strategy regions all values of the NPV are positive, and in the other region all are negative.

Table 3
Crucial Combinations of Two Input Variables

	Net Present Values at Usage Rates from 400 to 500				
Discount Rate	500	475	450	425	400
0.10	$16,814	$12,987	$9,161	$5,335	$1,509
0.11	$14,923	$11,200	$7,476	$3,753	$29
0.12	$13,082	$9,459	$5,835	$2,212	($1,411)
0.13	$11,288	$7,763	$4,238	$713	($2,812)
0.14	$9,541	$6,112	$2,683	($746)	($4,175)
0.15	$7,840	$4,505	$1,170	($2,165)	($5,500)
0.16	$6,185	$2,941	($302)	($3,545)	($6,788)
0.17	$4,573	$1,420	($1,734)	($4,887)	($8,040)
0.18	$3,004	($61)	($3,127)	($6,192)	($9,257)
0.19	$1,478	($1,501)	($4,481)	($7,460)	($10,440)
0.20	($6)	($2,902)	($5,798)	($8,693)	($11,589)

If most of the sensitivity in the model results from only two key variables, then this sort of analysis is very instructive. Even if there are more than two key variables, pair-wise analysis takes us at least one step closer to understanding the workings of the model in a realistic setting.

Graphic Analysis of Sensitivity

So far we have talked of interpreting tables to ascertain sensitivity. In practice, one is more likely to use graphs, so that the interaction between the input variable and the outcome is visible over a reasonable range of values for the input variable. One should keep in mind that sensitivity analysis is exploratory, not definitive, and therefore that making visible the patterns in the data is the first priority.

Sensitivity Curves

One simple and useful type of graph shows changes in net present values against changes in the input variable. It is easy to read switching values from such graphs, which provide a quick visual picture of how sensitive the outcome is to changes in each variable. If the change in the variable is presented on the graph as a percentage change (and thereby "standardized"), then it is possible to put the curves for two or more variables (calculated one at a time, of course) on the same graph. This is useful because the relative sensitivity of each input variable is indicated by the slopes of the curves. The more the NPV changes for a given change in each variable, the more sensitive it is to that variable, volatility being equal.

If one places NPV on the *X*-axis and the percentage change in the input variable on the *Y*-axis, then flat curves indicate a strong sensitivity. As one can see from the "sensitivity to benefit variables" shown in Figure 1, the NPV of the example project is more sensitive to variable B than it is to variable A. A small change in variable B leads to a relatively large change in the NPV. We can also see that the switching value of variable B is about 20% below its baseline (most likely) value, whereas the switching value of variable A is about 50% below its baseline value. Therefore we would give more attention to variable B than to variable A, again assuming both variables are equally volatile.

Figure 1
Sensitivity of the NPV to Benefit and Cost Variables

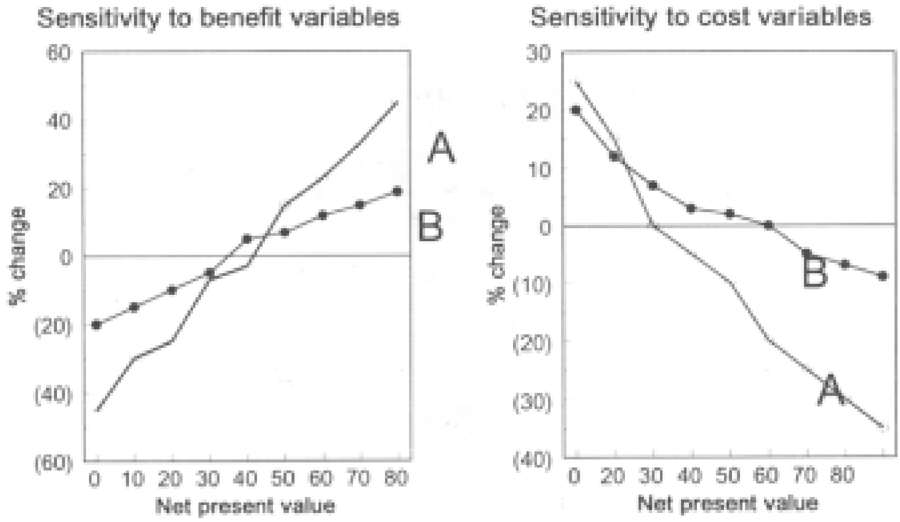

Spider Plots

Sensitivity plots can be consolidated to show many input variables on one chart. This type of consolidated chart is called a *spider plot*. The center of the spider plot is the NPV when all of the variables are at their baseline values. The lines of the spider chart show how the NPV changes as the values of each variable change, all others being held equal.

It is important to note that the lengths of the spider lines vary because each variable has its own plausible range within which it can change. The values of one variable might vary by only 10% up or down from its baseline value, whereas another variable might be highly uncertain, varying between plus 170% and minus 60% of its baseline value.

In Figure 2, the box that has been superimposed on the spider plot is an aid to interpretation. The top and bottom of the box indicate the level of a +5% and –5% change in the input variable. Wherever a tentacle of the spider plot crosses either the top or the bottom of the box, it shows what the net present value is, assuming a 5% change

Figure 2
Spider Plot—Sensitivity to Multiple Variables

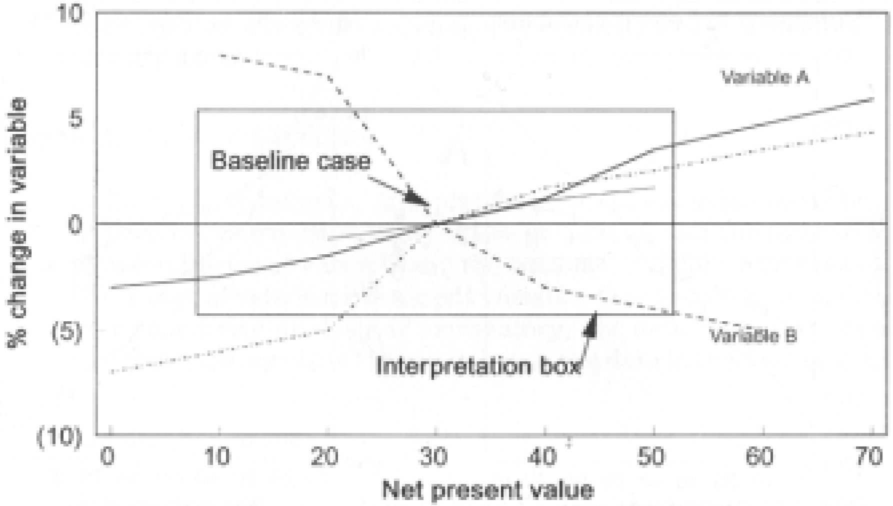

in the particular input variable. The left side of the box is at the zero NPV level, and therefore the tentacles cross this side at the switching value of the variable. The right side of the box is set at a net present value of 60 for symmetry (the baseline case is 30). These dimensions of the interpretation box are reasonable for this particular analysis, but will vary in other analyses—the box is a useful aid to reading the spider plot, nothing more. The spider plot describes how much each variable would have to change, other things held constant, to produce an NPV of zero (the switching ratio).

Tornado Charts

Tornado charts give us another quick, although partial, picture of relative sensitivity. Each bar in the tornado chart shows the range of the NPV if the variable in question is allowed to change from its highest to its lowest value. One can see from a quick glance at the shape of Figure 3, which has variables arranged in descending order from top to bottom, why it is called a tornado chart. Of course one should remember that not only the potential range of the variable, but also the probability that the value of the variable will move within that range (its *volatility*), is important to sensitivity.

Figure 3
A Tornado Diagram

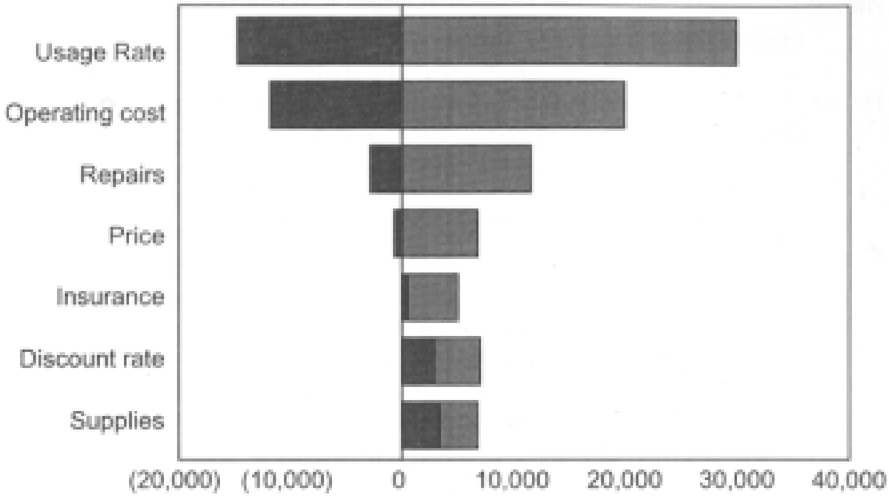

Action on Sensitivities

Once the key sensitivities have been identified among the risk variables, one by one, everything else held constant, one can start to think about managing risk.

- Are there input variables in the model that are correlated and therefore either dampen or enhance the influence each might have in isolation?
- Can diversification help? Are there other programs where the same variable works in the opposite direction?
- Can the value of the key variable be identified with more certainty by gathering more information, and if so, is it worth the cost to gather?

Once these questions are answered, the evaluator can formulate an action plan to minimize uncertainty and thereby to limit risk.

REFERENCES

Baird, B. (1989). Sensitivity analysis. Chapter 11 of *Managerial decisions under uncertainty* (pp. 352–375). New York: Wiley.

Hertz, D., & Thomas, H. (1984). *Practical risk analysis: An approach through case histories*. New York: Wiley.