

Javelin Diagrams: A Graphical Tool for Probabilistic Sensitivity Analysis

James C. Felli

Naval Postgraduate School, Defense Resources Management Institute (DRMI), Monterey, California 93943, jcfelli@nps.navy.mil

Gordon B. Hazen

Northwestern University, Evanston, Illinois 60208, hazen@iems.northwestern.edu

To demonstrate post hoc robustness of decision problems to parameter estimates, analysts may conduct a *probabilistic sensitivity analysis*, assigning distributions to uncertain parameters and computing the probability of decision change. In contrast to classical threshold proximity methods of sensitivity analysis, no appealing graphical methods are available to present the results of a probabilistic sensitivity analysis. Here we introduce an analog of tornado diagrams for probabilistic sensitivity analysis, which we call *javelin diagrams*. Javelin diagrams are graphical augmentations of tornado diagrams displaying both the probability of decision change and the information value associated with individual parameters or parameter sets. We construct javelin diagrams for simple problems, discuss their properties, and illustrate their realistic application via a probabilistic sensitivity analysis of a seven-parameter decision analysis from the medical literature.

Key words: decision analysis; information value; probabilistic sensitivity analysis; tornado diagrams; javelin diagrams

History: Received March 21, 2002. Accepted by Robert T. Clemen and Don N. Kleinmuntz on February 28, 2003, after 2 revisions.

1. Introduction

Sensitivity analysis today is a crucial element in any practical decision analysis, and can play any of several roles in the decision analysis process. In the *basis development* phase of modeling (Howard 1983, 1988), sensitivity analysis can be used to guide model development: If decisions are insensitive to changes in some aspect of the model, then there is no need to model that particular aspect in more detail (Howard 1983; Watson and Buede 1987, Chapter 7). For instance, Clemen (1996, Chapter 5) illustrates how *tornado diagrams* (e.g., Howard 1988) can be used to determine which deterministic variables have sufficient impact to be worth modeling probabilistically. Sensitivity analysis can also be used to give analysts *insight* into decision problems by identifying key variables (von Winterfeldt and Edwards 1986, Chapter 11; French 1986, pp. 342, 346). Sensitivity analysis may also be used to reduce the burden of assessing probabilities or utilities. For example, if a *strategy-region diagram* (e.g., Clemen 1996, Chapter 5) shows that the optimal decision remains so in a wide region of prob-

abilities or utilities, then no detailed assessment is required. Examples are given by Clemen (1996, Chapter 5), Keeney and Raiffa (1976, pp. 100, 203), Watson and Buede (1987, p. 270), and von Winterfeldt and Edwards (1986, Chapter 11). Finally, in what Howard (1983) calls the *defensible stage*, sensitivity analysis may be used in a post hoc fashion (that is, after the analysis is complete) to demonstrate to supportive or skeptical audiences the *robustness* of the analysis, or to point out that a decision is a close call (von Winterfeldt and Edwards 1986, p. 401; French 1986, p. 252; Keeney and Raiffa 1976, p. 460).

Analysts have long recognized the dimensionality limitations of graphically based sensitivity analysis in portraying the robustness of a decision analysis to variations in underlying parameter estimates. If graphical methods allow at most two- or three-way sensitivity analyses, how can one be sure that a decision analysis is robust to the simultaneous variation of 20 to 30 parameters? Probabilistic sensitivity analysis was introduced to address this issue. In a probabilistic sensitivity analysis, the analyst assigns

distributions to uncertain parameters and, as a measure of robustness, can thereby compute the probability of a change in the optimal alternative due to variation in arbitrarily many parameters.

In contrast to conventional methods of sensitivity analysis, which emphasize intuitively appealing graphical displays, the probabilistic sensitivity analysis approach is relatively devoid of graphical features. We introduce what might be termed the analog of tornado diagrams for probabilistic sensitivity analysis, which we call javelin diagrams. In a graphical display much like a tornado diagram, a javelin diagram displays not only the range of potential improvement offered by competing alternatives, but also the *probability of decision change* due to parametric variation and the *information value* associated with that variation. Although they are probabilistic analogs of tornado diagrams, javelin diagrams are intended to address post hoc robustness issues, whereas tornado diagrams were designed for use in the basis formulation stage of a decision analysis.

Probabilistic sensitivity analysis was first adopted in medical decision analyses (Doubilet et al. 1985, Critchfield and Willard 1986, Critchfield et al. 1986). It has remained popular in medicine and health economics (Manning et al. 1996, Sisk et al. 1997, Goodman et al. 1999, Lord and Asante 1999, Ng et al. 1999, Pasta et al. 1999, Murakami and Ohashi 2001, Williams et al. 2001), and has seen application in civil and environmental engineering (Cawlfeld and Wu 1993, Piggott and Cawlfeld 1996, Hamed 1999, Lin et al. 1999). Because probabilistic sensitivity analysis requires that distributions be assigned to all parameters, there may be an additional assessment burden imposed; however, for conventional sensitivity analysis purposes, a responsible analysis will already have specified a best estimate and plausible range for each parameter. From there it is a mechanical procedure to fit an appropriately chosen distribution with, for example, 95% of its probability mass in the plausible range and mean or mode equal to the best estimate, in which case no additional probability assessment is needed.

Conventional methods for sensitivity analysis, such as strategy-region diagrams or tornado diagrams, all rely implicitly on a *threshold proximity* view of sensitivity in which the analyst forms an intuitive judgment

of sensitivity by examining the proximity of a parameter's base value to the nearest threshold of decision change. A similar viewpoint is prevalent in probabilistic sensitivity analysis, where the analyst computes the probability of a threshold crossing. In recent work (Felli and Hazen 1998, 1999), we have advocated what might be termed a *value-focused* view (Keeney 1992) for probabilistic sensitivity analysis, in which the analyst assesses sensitivity by computing the *value of information* for a parameter or parameter set. Empirical results indicate that in comparison with information value, threshold proximity approaches tend to significantly overestimate problem sensitivity (Felli and Hazen 1998, 1999).

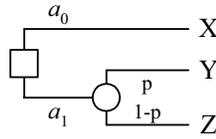
All measures of post hoc robustness attempt to quantify the degree to which parameter uncertainty may produce a change in the optimal choice. That is the point of view we adopt in this paper. Under this perspective, the sensitivity of the optimal *payoff* to parameter variation is unimportant unless it produces a change in the optimal *choice*. For example, if parameter variation cannot produce a change in optimal choice, then threshold proximity is nonexistent, the probability of decision change is zero, and the information value is also zero, *regardless* of how much the optimal payoff changes with parameter variation.

This paper is organized as follows. In §2, we review tornado diagrams and introduce necessary terminology. In §3, we move to probabilistic sensitivity analysis and introduce javelin diagrams. Section 4 addresses the general multiparameter, multiple-alternative version of javelin diagrams and extends the discussion to incorporate nonneutral risk attitudes. In §5, we apply the tool to a previously published medical decision analysis. We then close with a discussion of the relative advantages of javelin diagrams.

2. Tornado Diagrams

Consider the decision problem in Figure 1 in which the probability p and the payoffs X, Y, Z are parameters, and the decision maker wishes to maximize expected payoff. Let $\Pi = (p, X, Y, Z)$ be the set of problem parameters, and let $\Pi_0 = (p_0, X_0, Y_0, Z_0)$ be the set of *base parameter values*, the estimated values of those parameters. We will denote the expected payoff

Figure 1 A Simple Decision Problem with Parameters $p, X, Y,$ and Z



of an alternative a as a function of its parameters Π by $E[V_a | \Pi]$. In this case, we have

$$E[V_{a_0} | \Pi] = X,$$

$$E[V_{a_1} | \Pi] = pY + (1 - p)Z.$$

The alternative that maximizes expected payoff is a function $a_0(\Pi)$ of problem parameters. We let $a_0 = a_0(\Pi_0)$ be the optimal alternative given $\Pi = \Pi_0$ and will refer to a_0 as the *base-optimal alternative* (BOA).

Base values and plausible ranges for the four parameters in this problem are shown in Table 1. Using these data, we see that $E[V_{a_0} | \Pi_0] = \$25$ and $E[V_{a_1} | \Pi_0] = -\$20$, so a_0 is the BOA.

For the purposes of this illustration, we assume that successive cycles of model refinement are complete, and the analyst wishes to conduct a post hoc robustness analysis on the parameters $p, X, Y,$ and Z . A tornado diagram that accomplishes this is given in Figure 2. The tornado diagram depicts the range of possible payoff gains,

$$E[\Delta V_{a_1} | \xi, \Pi_0 \setminus \xi_0] = E[V_{a_1} | \xi, \Pi_0 \setminus \xi_0] - E[V_{a_0} | \xi, \Pi_0 \setminus \xi_0],$$

obtainable by varying the parameter ξ over its plausible range, for each of the four parameters $\xi \in \Pi$. Here, $\Delta V_{a_1} = V_{a_1} - V_{a_0}$, and the notation $\Pi_0 \setminus \xi_0$ indicates the set Π_0 of base-value parameters with ξ_0 excluded. For example, when $\xi = p$, we have

$$E[\Delta V_{a_1} | p, \Pi_0 \setminus p_0] = E[V_{a_1} | p, \Pi_0 \setminus p_0] - E[V_{a_0} | p, \Pi_0 \setminus p_0]$$

$$= (pY_0 + (1 - p)Z_0) - X_0$$

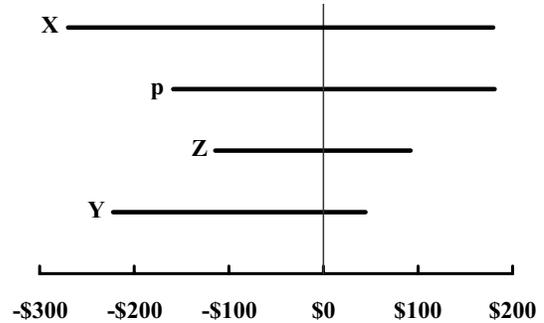
$$= p(-133.33) + (1 - p)(206.67) - 25$$

$$= 181.67 - 340p.$$

Table 1 Base Parameter Values and Ranges for the Decision Problem in Figure 1

Parameter	Base	Minimum	Maximum
p	0.667	0	1
X (\$)	25	-200	250
Y (\$)	-133.33	-400	0
Z (\$)	206.67	0	620

Figure 2 A Tornado Diagram in Which Deviations to the Right of Zero Indicate Potential Gains Over the Payoff of the Base Optimal Alternative a_0



The range of this function of p is depicted as the p -bar in Figure 2. The other three payoff gain functions are

$$E[\Delta V_{a_1} | X, \Pi_0 \setminus X_0] = (p_0 Y_0 + (1 - p_0) Z_0) - X$$

$$= -20 - X,$$

$$E[\Delta V_{a_1} | Y, \Pi_0 \setminus Y_0] = (p_0 Y + (1 - p_0) Z_0) - X_0$$

$$= 43.82 + 0.667Y,$$

$$E[\Delta V_{a_1} | Z, \Pi_0 \setminus Z_0] = (p_0 Y_0 + (1 - p_0) Z) - X_0$$

$$= -113.93 + 0.333Z.$$

The quantity $E[\Delta V_{a_1} | \xi, \Pi_0 \setminus \xi_0]$ is the payoff gain the decision maker could expect from a_1 over a_0 as a function of ξ . Using this formulation, zero becomes the natural reference point: The positive portion of a tornado bar for ξ corresponds to values of ξ for which the competing alternative is superior to the BOA. It is also possible to produce a tornado diagram that depicts the range of the alternative- a_1 payoffs $E[V_{a_1} | \xi, \Pi_0 \setminus \xi_0]$ as a function of ξ , with reference point the alternative- a_0 payoff $E[V_{a_0} | \Pi_0]$ (e.g., see the Eagle Airlines example in Clemen 1996, Chapter 5). This format has the disadvantage that sensitivity to ξ is not properly depicted when $E[V_{a_0} | \Pi]$ is also a function of ξ . For our example, either version of the tornado diagram would be adequate.

3. Javelin Diagrams

An alternate method of checking post hoc robustness is to perform a probabilistic sensitivity analysis. It is easy to illustrate probabilistic sensitivity analysis using tornado diagrams. Suppose, for example,

that the parameters $\Pi = (p, X, Y, Z)$ from our previous example are assigned beta densities as described in Table 2. These densities are depicted graphically in Figure 3. We assume no probabilistic dependence between the parameters.

Using these parameter densities, we can determine the distribution of each of the payoff gains $E[\Delta V_{a_1} | \xi, \Pi_0 \setminus \xi]$ in the tornado diagram we created in Figure 2. In simple cases, these distributions may be determined algebraically; in more complex cases, numerical methods such as Monte Carlo simulation may be required. We used Monte Carlo simulation to generate the payoff-gain densities for the decision problem of Figure 1. These are presented in Figure 4.

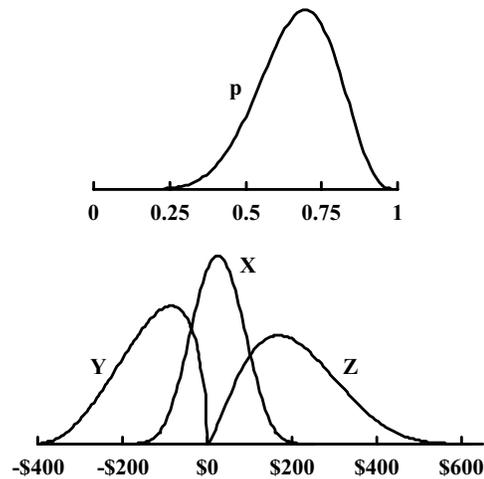
Another indicator of sensitivity in tornado diagrams is the absolute length of the tornado bar to the right of zero, which equals the *maximum payoff gain* due to variation in the corresponding parameter. For example, in Figure 4, the maximum payoff gain due to variation in X is \$179.89. The maximum payoff gain for Y is \$43.82. (In line with our discussion in the introduction, it is important to notice we should not be interested in the *minimum* payoff gain because the BOA remains optimal in the region to the left of zero.)

A more revealing sensitivity indicator would be the *average positive payoff gain*, equal to

$$E[E[\Delta V_{a_1} | \xi, \Pi_0 \setminus \xi]^+],$$

and formed by using the parameter distribution to average the positive parts of the payoff gains. We do not average the payoff gain itself because parameter variation producing a negative payoff gain does not result in suboptimality for the BOA, and therefore produces no decision change. We will use the more succinct term—*expected improvement*—to refer to average positive payoff gain. For many common decision

Figure 3 Parameter Densities for the Decision Problem of Figure 1



problems (including the decision problem of Figure 1), the expected improvement for a parameter ξ is equal (see the Appendix) to the *expected value of perfect information on ξ* , which we denote $EVPI_{\xi}$.

As we have noted, in the decision problem of Figure 1, we have

$$E[\Delta V_{a_1} | \Pi] = pY + (1 - p)Z - X.$$

Therefore, the expected improvement for p , for example, is given by

$$E[E[\Delta V_{a_1} | p, \Pi_0 \setminus p_0]^+] = E[(pY_0 + (1 - p)Z_0 - X_0)^+].$$

Expected improvement is not available directly in the payoff-gain density diagrams of Figure 4, although it

Figure 4 Payoff-Gain Densities for the Decision Problem of Figure 1, Imposed Over the Tornado Bars of Figure 2

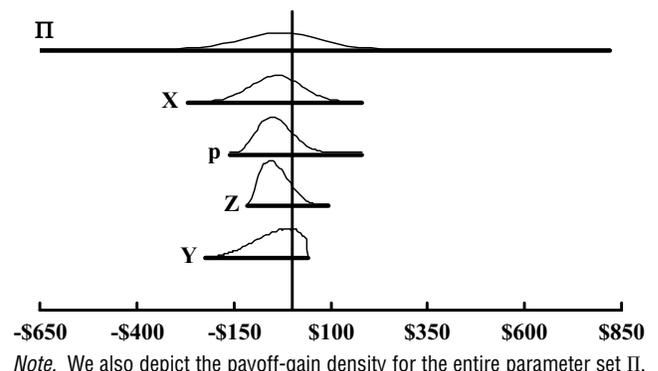


Table 2 Base Values, Ranges, and Distributions for the Parameters of the Decision Problem in Figure 1

Parameter	Base	Minimum	Maximum	Beta(α, β) densities	
				α	β
p	0.667	0	1	9	4.5
X	\$25.00	-\$200	\$250	7	7
Y	-\$133.33	-\$400	\$0	3.15	1.575
Z	\$206.67	\$0	\$620	2.475	4.95

Note. All parameters are distributed as scaled beta densities with base values equal to expected values.

Note. We also depict the payoff-gain density for the entire parameter set Π .

may be calculated by multiplying the probability of decision change by the conditional expected payoff given a decision change (i.e., the center of gravity of the portion of the payoff-gain density to the right of zero).

For javelin diagrams as we now define them, both the expected improvement and likelihood of improvement may be viewed graphically. In a javelin diagram, instead of the density of each payoff gain $E[\Delta V_{a_1} | \xi, \Pi_0 \setminus \xi]$, we graph the complementary cumulative distribution of this quantity; that is, we graph the function

$$G_\xi(\Delta v) = P(E[\Delta V_{a_1} | \xi, \Pi_0 \setminus \xi] > \Delta v)$$

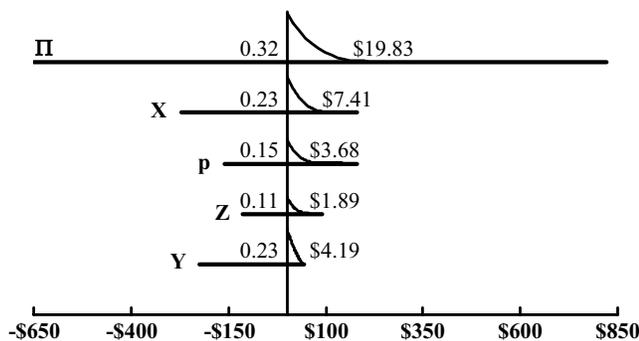
for $\Delta v \geq 0$. The javelin diagram consequent to Figure 4 is shown in Figure 5. The diagram gets its name from the graphical shaft formed by plotting the range of payoff gain and the head formed by the graph of $G_\xi(\Delta v)$ to the right of $\Delta v = 0$. The likelihood of improvement $G_\xi(0)$ is the height of the javelin head and the expected improvement is the area to the right of zero under the curve $G_\xi(\Delta v)$. In Figure 5, the number to the left of the vertical line at zero for each ξ -bar is the probability $G_\xi(0)$ that the BOA loses optimality due to variation in ξ . The number to the right of the zero line is the expected improvement, equal to $EVPI_\xi$ in this case.

The simultaneous display of expected improvement and likelihood of improvement is a key feature of

javelin diagrams. The javelin diagram in Figure 5 shows, for example, that (1) there is a 32% chance that the competing alternative will outperform the BOA due to simultaneous variation in all parameters, and (2) by observing all parameters and adapting his choice as necessary, the decision maker could increase his payoff by an average of \$19.83 over what he would expect from the BOA (i.e., $EVPI_\Pi = \$19.83$). We contend that expected improvement is a better sensitivity indicator than probability of decision change because it measures the *degree of impact* of a decision change rather than merely indicating its likelihood, and measures this impact in units the decision maker cares about—*payoff gains*. The javelin diagram provides intuitive support for this point of view. Compare the javelin for X in Figure 5 with the javelin for Y: The larger javelin head for X gives the intuitive—and we argue, correct—impression that the problem is more sensitive to X than to Y, even though both parameters produce the same 23% likelihood of decision change.

How should one use the information provided in a javelin diagram? In this situation, the decision maker might reason as follows: In a problem in which the base-optimal payoff is \$25, an expected improvement of $\$19.83 = EVPI_\Pi$ seems substantial. Therefore, the problem seems jointly sensitive to all its parameters. Payoff gains of \$3 or greater seem significant, so the problem appears sensitive to the individual parameters p, X, Y as each has an EVPI value in excess of \$3. However, the problem may not be sensitive to the parameter Z individually, as its EVPI is only \$1.89. We state these assertions tentatively because they depend on a subjective declaration by the decision maker as to what constitutes a significant improvement in payoff.

Figure 5 A Javelin Diagram for the Decision Problem of Figure 1



Notes. The javelin curves indicate the probability that parametric variation yields payoff gain at least as great as that noted on the horizontal axis. Probabilities displayed to the left of each javelin head are the javelin height, equal to the probability of decision change. The values displayed to the right of each javelin are the areas under the javelin curves, equal to the expected value of perfect information of that parameter or parameter set.

4. Multiple Alternatives and Nonneutral Risk Attitude

Multiple Alternatives

Both tornado diagrams and javelin diagrams may be readily adapted to the case of multiple alternatives. Suppose a_0 is the BOA and let $\Delta V_a = V_a - V_{a_0}$ be the difference between the payoff V_a under a and the payoff V_{a_0} under a_0 . If ξ is a parameter or set of parameters, let $a_0(\xi)$ be the optimal alternative as a function of ξ when all other parameters $\Pi \setminus \xi$ lie at their base

values. Here alternatives are chosen from a set A of feasible actions. Then, the range of the payoff-gain function $\xi \rightarrow E[\Delta V_{a_0(\xi)} | \xi, \Pi_0 \setminus \xi_0]$ may be graphed in a tornado diagram. Here, the payoff gain $\Delta V_{a_0(\xi)} = V_{a_0(\xi)} - V_{a_0}$ is always nonnegative, so this tornado diagram would be truncated to the left at zero. To obtain an untruncated version, let $a'_0(\xi)$ be the optimal alternative as a function of ξ in the reduced feasible set $A \setminus a_0$. The payoff gain $E[\Delta V_{a'_0(\xi)} | \xi, \Pi_0 \setminus \xi_0]$ may be positive or negative. The graph of the range of the function $\xi \rightarrow E[\Delta V_{a'_0(\xi)} | \xi, \Pi_0 \setminus \xi_0]$ will yield an untruncated tornado diagram for payoff gains. Because $E[\Delta V_{a_0(\xi)} | \xi, \Pi_0 \setminus \xi_0]$ is equal to $E[\Delta V_{a'_0(\xi)} | \xi, \Pi_0 \setminus \xi_0]$ when the latter is nonnegative, the truncated and untruncated tornado diagrams will be identical to the right of zero.

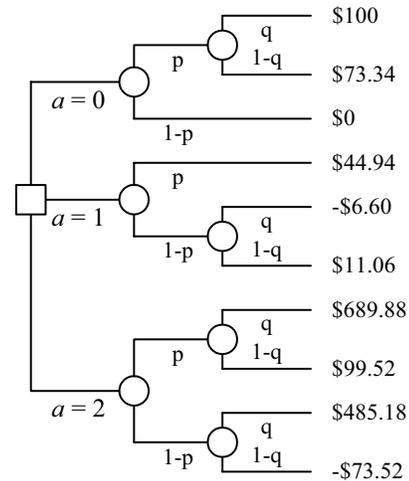
Given a probability distribution for the parameter ξ , one can form a javelin diagram by first calculating the complementary cumulative distribution $G_\xi(\Delta v) = P(E[\Delta V_{a_0(\xi)} | \xi, \Pi_0 \setminus \xi_0] > \Delta v)$ and then placing the graph of $G_\xi(\Delta v)$ for $\Delta v \geq 0$ over the corresponding untruncated tornado bar for the parameter ξ . The expected improvement $E_\xi[E[\Delta V_{a_0(\xi)} | \xi, \Pi_0 \setminus \xi_0]]$ is the area to the right of zero under $G_\xi(\Delta v)$. The probability of decision change due to uncertainty in ξ is $G_\xi(0)$. As before, the expected improvement will often be equal to the information value of ξ (see Appendix).

Example: The Nonequivalence of Javelin and Tornado Diagrams

The javelin diagram is the natural extension of the tornado diagram for the purpose of incorporating prior probability distributions over parameters. Consequently, one might conjecture that when all parameters are uniformly distributed over their ranges, the javelin diagram conveys the same information and leads to the same sensitivity conclusions as the tornado diagram. The following example shows that this conjecture is false for three or more alternatives.

Consider the three-alternative example in Figure 6. There are two parameters p, q with base values p_0, q_0 , and ranges given in Table 3. At the base values, the alternatives $a = 0, 1, 2$ have expected payoffs: $E[V_0 | p_0, q_0] = \$19.48$, $E[V_1 | p_0, q_0] = \$19.08$, $E[V_2 | p_0, q_0] = \$5.48$. Therefore, the BOA is $a_0 = 0$.

Figure 6 A Hypothetical Three-Alternative Decision Problem with Two Probability Parameters p and q



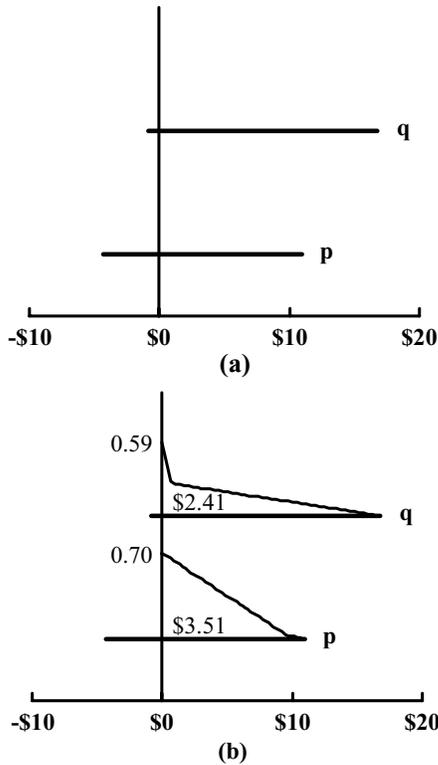
Using the ranges and base values for p and q in Table 3, we can construct the untruncated tornado diagram for payoff gains in Figure 7a. If we assign distributions to p and q that are *uniform* over the ranges given in Table 3, we can construct the javelin diagram in Figure 7b.

The tornado diagram indicates that the problem is sensitive to both parameters, but more so to q than to p . The javelin diagram, on the other hand, indicates unambiguously that the problem is more sensitive to parameter p than to parameter q , exactly the reverse of the conclusion from the tornado diagram. Variation in p has a 70% chance of changing the optimal decision, with corresponding information value $EVPI_p = \$3.51$. The problem is also sensitive to q but less so, with the probability of decision change equal to 54% and information value $EVPI_q = \$2.40$. The range information in the tornado diagram does not adequately account for sensitivity in this case, even though both parameters were assigned uniform distributions. *Even with all parameters distributed uniformly, a javelin dia-*

Table 3 Parameter Base Values and Ranges for the Decision Problem in Figure 6

Parameter	Base	Minimum	Maximum
p	0.26	0.01	0.51
q	0.06	0.005	0.115

Figure 7 (a) Tornado Diagram for the Decision Problem in Figure 6
 (b) Javelin Diagram for the Decision Problem in Figure 6



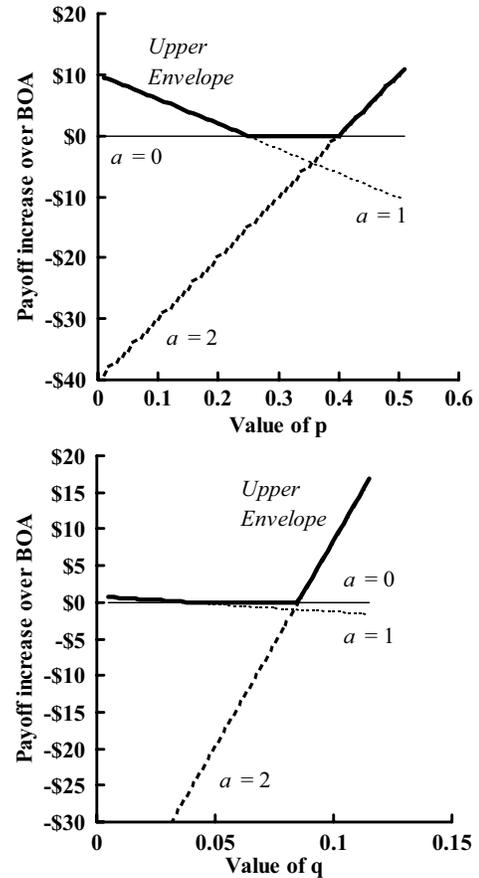
Notes. (a) The bars represent the payoff gains over the BOA as p and q vary independently over their ranges. (b) Numbers to the left denote the probability that a_0 loses optimality due to parametric variation; numbers to the right, the expected value of perfect parameter information for the parameter. The javelin curves indicate the probability that parametric variation yields payoff gain at least as great as on the horizontal axis.

gram need not convey the same sensitivity conclusions as a tornado diagram.¹

The reader may get some feeling for how this happens by examining Figure 8, in which graphs of the functions $\xi \rightarrow E[\Delta V_a | \xi, \Pi_0 \setminus \xi_0]$ appear for $\xi = p, q$ and $a = 0, 1, 2$ in the decision problem of Figure 6. The upper envelope in these graphs is the payoff-gain function $\xi \rightarrow E[\Delta V_{a_0(\xi)} | \xi, \Pi_0 \setminus \xi_0]$. The projection of each upper envelope onto the vertical axis (the vertical range) constitutes the nonnegative por-

¹ The exception is when there are only two alternatives. When there are only two alternatives and payoffs are multilinear in parameters ξ , we can show that sensitivity conclusions using a tornado diagram will always match those using a javelin diagram with independent, uniformly distributed parameters having means equal to base values. This is beyond the scope of this paper.

Figure 8 One-Way Sensitivity Analysis Graphs of the Functions $\xi \rightarrow E[\Delta V_a | \xi, \Pi_0 \setminus \xi_0]$ for $\xi = p, \xi = q$, and $a = 0, 1, 2$ in the Decision Problem of Figure 6



Note. A tornado diagram depicts the vertical range of the upper envelope, which is larger for q than for p . However, because the parameters are distributed uniformly, information value is equal to the average height of the upper envelope and the probability of decision change is the proportion of the horizontal axis where the upper envelope exceeds zero. Both are greater for p than for q .

tion of the corresponding tornado bar. On the other hand, the average height of the upper envelope is $E_\xi[E[\Delta V_{a_0(\xi)} | \xi, \Pi_0 \setminus \xi_0]]$ and is equal to the information value $EVPI_\xi$. This quantity is larger for $\xi = p$ than for $\xi = q$, even though the vertical range is greater for $\xi = q$.

Nonneutral Risk Attitude

We have introduced the javelin diagram as a sensitivity analysis tool for the cases in which payoff is monetary and the decision maker is risk neutral. In fact, no such restrictions are necessary: The

notion of information value as expected improvement and the corresponding javelin display can be extended to the case of nonneutral risk attitude and even to the case in which there is no underlying monetary attribute. Because the approach we take is unorthodox—perhaps even heterodox to some—we present it here in detail.

Suppose a utility function $u(x)$ has been assessed over problem outcomes x , and let U_a be a random variable equal to the utility achieved under alternative a . Let a^* be the alternative maximizing expected utility $E[U_a]$, and let $a^*(\xi)$ be the alternative a maximizing the conditional expected utility $E[U_a | \xi]$ given the value of parameter ξ . Define

$$EU_\xi = E_\xi[E[U_{a^*(\xi)} | \xi]],$$

the expected utility of learning the value of ξ . Similarly, let $EU_\emptyset = E[U_{a^*}]$ be the optimal expected utility under no further information.

Larger values of EU_ξ indicate greater problem sensitivity to ξ , so the quantities EU_ξ may be used to rank problem sensitivity to different parameter sets ξ . This is the first role of any sensitivity measure. The analyst may find it helpful to form the *expected improvement in utility*,

$$EUI_\xi = EU_\xi - EU_\emptyset.$$

However, this measure is merely for convenience. Although it is possible (see below) and even desirable to assign an information theoretic meaning to EUI_ξ , there is no need to do so for the purposes of sensitivity analysis.

We require a second role of any sensitivity measure, namely the declaration (yes/no) of problem sensitivity to a set ξ of parameters. To this end, let \tilde{x}^* be the random outcome corresponding to the optimal alternative a^* . We require the decision maker or analyst to declare *an outcome having minimum significant improvement*, that is, a (possibly random) outcome \tilde{x}^+ having the properties that (1) \tilde{x}^+ is preferred to \tilde{x}^* , (2) \tilde{x}^+ is *significantly* preferred to \tilde{x}^* , and (3) no other random or nonrandom outcome \tilde{x} that is less preferred than \tilde{x}^+ can be significantly preferred to \tilde{x}^* . Declaring an outcome \tilde{x}^+ having minimum significant improvement over \tilde{x}^* is a completely subjective judgment by the analyst or decision maker, and is not a property of

the assessed utility function u . Nevertheless, the outcome \tilde{x}^+ having minimum significant improvement over \tilde{x}^* induces a *minimum significant increase* Δu^+ in utility, given by

$$\Delta u^+ = E[u(\tilde{x}^+)] - E[u(\tilde{x}^*)].$$

The value Δu^+ of the minimum significant increase in utility is intrinsic to the particular decision problem at hand, as it depends on how the utility function is scaled and the optimal random outcome \tilde{x}^* . It is not possible to declare a minimum significant improvement in utility valid across all decision problems. Nevertheless, for the problem at hand, no outcome \tilde{x} whose improvement in utility is less than Δu^+ will constitute a significant improvement over \tilde{x}^* .

We propose to extend the notion of significant improvement to information sources, and say that perfect information about a parameter set ξ *constitutes a significant improvement* if the expected utility EU_ξ of learning ξ exceeds the expected utility $E[u(\tilde{x}^+)]$ of \tilde{x}^+ , or equivalently, if the expected improvement in utility EUI_ξ by learning ξ exceeds the minimum significant utility improvement Δu^+ . We propose to declare a decision problem sensitive to a parameter set ξ if learning the value of ξ constitutes a significant improvement.

Although this approach to probabilistic sensitivity analysis has an information-theoretic basis, there is no requirement that the expected improvement in utility EUI_ξ itself has an information-theoretic meaning. Nevertheless, EUI_ξ can in fact be interpreted as a form of information value, as we shall now explain. Suppose we let x^- be an outcome (call it the *negative outcome*) having utility $u(x^-)$ strictly smaller than optimal expected utility EU_\emptyset . Define the *probability price* PP_ξ of the parameter set ξ to be the largest probability of the negative outcome x^- one is willing to accept to obtain perfect information about ξ . The probability price of ξ is a measure of information value, and satisfies

$$PP_\xi u(x^-) + (1 - PP_\xi) EU_\xi = EU_\emptyset.$$

Rearranging and using $EUI_\xi = EU_\xi - EU_\emptyset$, we have

$$\frac{PP_\xi}{1 - PP_\xi} = \frac{EUI_\xi}{EU_\emptyset - u(x^-)}.$$

We see, therefore, that the expected improvement in utility EUI_ξ is proportional to the odds price $PP_\xi/(1 - PP_\xi)$, that is, EUI_ξ is proportional to the largest odds of the negative outcome x^- one would accept to learn ξ . Rescaling utility if necessary, we can choose x^- so that $EU_\emptyset - u(x^-) = 1$, in which case EUI_ξ is the largest odds of a unit reduction of utility from EU_\emptyset that one would accept to learn ξ .

Probability price and its relationship to expected improvement in utility were discussed by Hazen and Souderpandian (1999). When outcomes are monetary and risk attitude is not neutral, the conventional definition is that information value $EVPI_\xi$ is the most one would be willing to pay to learn ξ , that is, $EVPI_\xi$ satisfies

$$E[u(x^* - EVPI_\xi) | \xi] = EU_\emptyset.$$

Hazen and Souderpandian (1999) show that when risk attitude is constant (i.e., a linear or exponential utility function), then expected improvement in utility EUI_ξ and conventionally defined information value $EVPI_\xi$ rank information sources identically, but otherwise, they may differ. The advantage to using EUI_ξ is that it is computationally and analytically simpler when risk attitude is not constant, and it generalizes to situations in which outcomes are not monetary.

In the preceding discussion, we have glossed over a practical detail involving base estimates for parameters. The optimal expected utility may be written as

$$EU_\emptyset = E[U_{a^*}] = E_{\Pi}[E[U_{a^*} | \Pi]].$$

In the common special case in which $E[U_{a^*} | \Pi]$ is a *multilinear* function of the parameters Π (i.e., linear in each component of Π) and parameters are independent, then we have

$$EU_\emptyset = E[U_{a^*} | \bar{\Pi}],$$

where $\bar{\Pi} = E[\Pi]$ is the mean of Π . If the analyst has taken the base-optimal level Π_0 of parameters Π to be the mean $\bar{\Pi}$, then it follows that a^* is the base-optimal alternative a_0 obtained by maximizing $E[U_{a^*} | \Pi_0]$. Similarly, for any parameter set ξ , $a^*(\xi)$ will equal $a_0(\xi)$. (See the Appendix for details.) However, if these conditions fail, that is, if parameters are not independent, if multilinearity does not hold, or if base parameter

levels are not equal to their mean values, then a^* may differ from a_0 , and $a^*(\xi)$ may differ from $a_0(\xi)$. When this occurs, and the analyst computes expected improvements based on a_0 and $a_0(\xi)$ instead of a^* and $a^*(\xi)$, then the information-theoretic basis for the procedure is only approximately maintained.

5. A Realistic Example

As an example of constructing and interpreting a javelin diagram, we will revisit the problem of management of suspected giant cell arteritis posed by Buchbinder and Detsky (1992). This problem illustrates the use of javelin diagrams on a realistic-size problem, and the use of expected improvement in utility as a sensitivity measure, as discussed in the previous section.

Buchbinder and Detsky (1992) considered four alternatives in this problem: (A) Treat None, (B) Biopsy and Treat Positive, (C) Biopsy and Treat All, and (D) Treat All. The structure we provide below is the same as Buchbinder and Detsky (1992) with the exception of the notation we adopted for ease of exposition. The four alternatives are presented in Figures 9a–9d.

Table 4 provides a summary of the parameter values used by Buchbinder and Detsky (1992) in their analysis and the range of values they used for a sensitivity analysis with the probability of having giant cell arteritis (GCA) set at 0.8. The reader may note that the expected utilities $E[U_a | \Pi]$ are all multilinear in the parameters Π . With all parameters held at base value, Biopsy and Treat Positive was optimal with an expected payoff of 0.837 on a utility scale from -0.325 (all disutilities present) to 1 (perfect health). The expected utilities of the four treatment alternatives at base parameter values are provided in Table 5. Note that all utilities in this analysis were physician assessed. The analysis therefore takes the perspective of the “benevolent physician.”

Figure 9a The Decision Tree for the Treat None Alternative

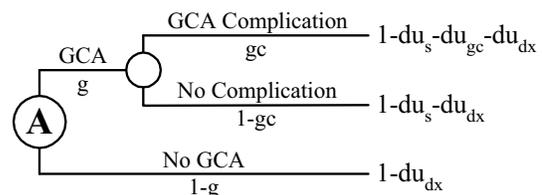
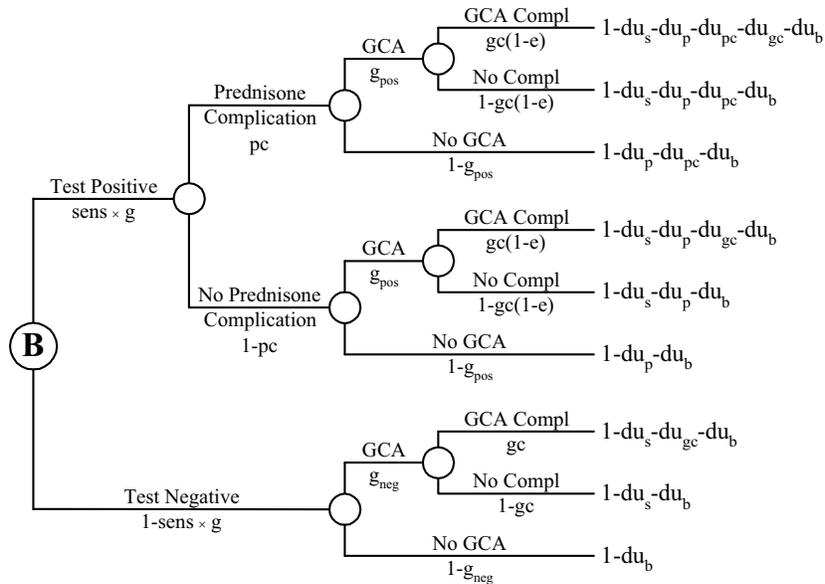
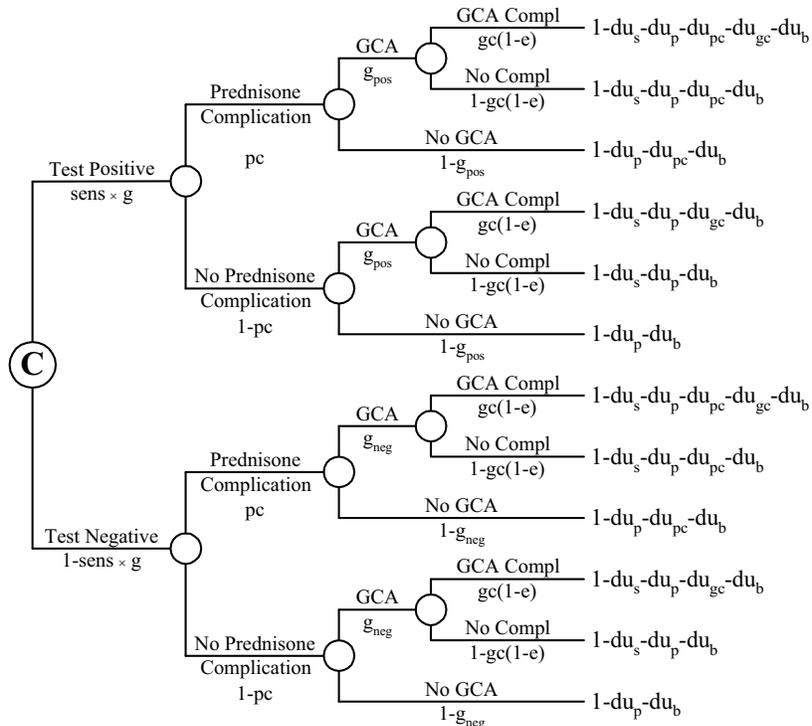


Figure 9b The Decision for the Biopsy and Treat Positive Alternative



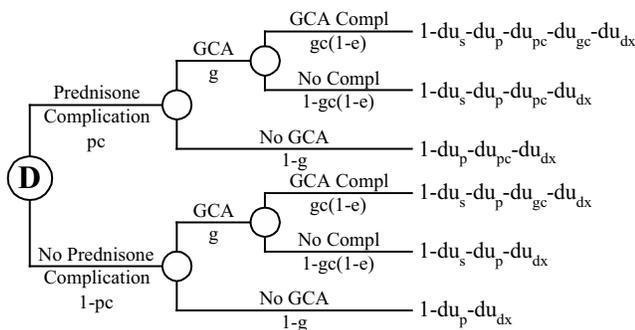
Notes. The probability of a positive test is equal to $sens \times g$. The terms g_{pos} and g_{neg} refer to the probability of having GCA given a positive or negative test result. Because the specificity of the test is 1, $g_{pos} = 1$ and $g_{neg} = (1 - sens)g / (1 - sens \times g)$.

Figure 9c The Decision Tree for the Biopsy and Treat All Alternative



Notes. The probability of a positive test is equal to $sens \times g$. The terms g_{pos} and g_{neg} refer to the probability of having GCA, given a positive or negative test result. Because the specificity of the test is 1, $g_{pos} = 1$ and $g_{neg} = (1 - sens)g / (1 - sens \times g)$.

Figure 9d The Decision Tree for the Treat All Alternative



The threshold column in Table 4 notes the value of the parameter required to cause the BOA (Biopsy and Treat Positive) to lose optimality when all other parameters remain fixed at their base values. Except for the beta density parameters, Buchbinder and Detsky (1992) provided the data in Table 4. We assumed parameters to be probabilistically independent and determined the beta density parameters by fitting beta distributions to Buchbinder and Detsky (1992) data, using the base values as mean values and capturing 95% of the probability mass in the range defined by the parameter’s minimum and maximum plausible values. We let the remaining 5% spill equally to either side of this range if possible (that is, if the upper plausible limit was less than one and the lower limit greater than zero—otherwise we filled the

plausible range with 97.5% or 100% probability mass, as appropriate).

The close proximity of threshold values to parameter base values implies that a small parameter value deviation away from base value could cause the BOA to lose optimality. Based on this, Buchbinder and Detsky (1992), in their post hoc robustness analysis, declared the BOA sensitive to the seven nonconstant parameters in Table 4 when $g = 0.8$.

To check these sensitivity conclusions, we constructed a tornado diagram and a javelin diagram. Our tornado diagram depicting utility improvements is provided in Figure 10. The optimality of the BOA not only appears to be clearly sensitive to variation in the sens parameter, but extremely susceptible to variation in du_p . The javelin diagram corresponding to the beta densities in Table 4 is provided in Figure 11a. The javelin heads were determined via Monte Carlo simulation. Figure 11b is a rescaled version of Figure 11a that gives a better view of the javelin heads.

Areas under the javelin heads are expected improvement *in utility*, discussed in the previous section. For this analysis, an expected improvement in utility of Δu for parameter ξ indicates that the possibility of learning ξ is worth a $\Delta u/(1 + \Delta u)$ chance of falling one unit in utility from the base-optimal level 0.837. A drop of one unit in utility from 0.837, that is, from 0.837 to -0.163 , amounts to a total utility decrement of $1 - (-0.163) = 1.163$ from the well state,

Table 4 The Data Used by Buchbinder and Detsky (1992) in Their Analysis

Parameters	Symbols	Base	Minimum	Maximum	Threshold	Beta(α, β)	
						α	β
P[having GCA]	g	0.8	—	—	—	—	—
P[developing severe complications of GCA]	gc	0.3	0.05	0.5	0.31	4.179	11.011
P[developing severe iatrogenic side effects]	pc	0.2	0.05	0.5	0.19	2.647	10.589
Efficacy of high dose prednisone	e	0.9	0.8	1	0.94	27.787	3.087
Sensitivity of temporal artery biopsy	$sens$	0.83	0.6	1	0.82	7.554	1.547
Specificity of temporal artery biopsy	$spec$	1	—	—	—	—	—
D(major complication from GCA)	du_{gc}	0.8	0.3	0.9	0.84	27.454	6.864
D(prednisone therapy)	du_p	0.08	0.03	0.2	0.08	4.555	52.380
D(major iatrogenic side effect)	du_{pc}	0.3	0.2	0.9	0.28	15.291	35.680
D(having symptoms of GCA)	du_s	0.12	—	—	—	—	—
D(having a temporal artery biopsy)	du_b	0.005	—	—	—	—	—
D(not knowing the true diagnosis)	du_{dx}	0.025	—	—	—	—	—

Notes. The minimum and maximum values depict each parameter’s range for sensitivity analysis. The threshold value indicates the parameter value at which the BOA lost optimality to another alternative. The values $P[\cdot]$ and $D(\cdot)$ designate the probability of an event and the disutility associated with an event. Buchbinder and Detsky (1992) assigned $du_s = 0$ when treatment was given because the effectiveness of therapy completely alleviated the patient’s symptoms.

Table 5 The Expected Utilities of the Four Alternative Treatments for Suspected Giant Cell Arteritis When All Parameters Are Set to Their Base Values in Table 4

Treatment alternative	EU at base values
Treat None	0.687
Biopsy and Treat Positive	0.837
Biopsy and Treat All	0.836
Treat All	0.816

Note. With all parameters at their base values, the Buchbinder and Detsky (1992) utility scale ranged from -0.325 (all disutilities present) to 1 (perfect health).

approximately equivalent here to the simultaneous occurrence of major complications from GCA and major iatrogenic side effects.

Figures 11a and b substantiate the results of Buchbinder and Detsky’s (1992) threshold analyses in that the probability of another alternative yielding a higher expected utility than the BOA is high. However, the all-parameter information value is a utility improvement of merely 0.0104, indicating that learning the true values of all parameters is worth taking at most a $0.0104/(1 + 0.0104) = 1.03\%$ chance at the simultaneous occurrence of major complications from GCA and major iatrogenic side effects. The significance of a utility improvement of 0.0104 over the base-optimal utility 0.837 is a subjective judgment. However, it is worth noting that the smallest utility decrement noted by Buchbinder and Detsky (1992) in their analysis was the disutility for having a temporal artery biopsy, equal to 0.005. If one could eliminate the disutility of temporal artery biopsy from the optimal policy Biopsy and Treat Positive, its expected

Figure 10 The Payoff-Gain Tornado Diagram for the Giant Cell Arteritis Problem

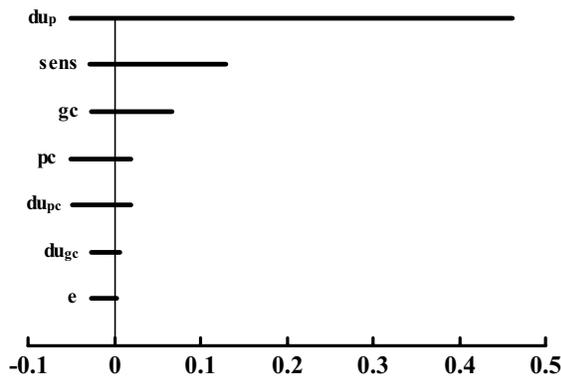
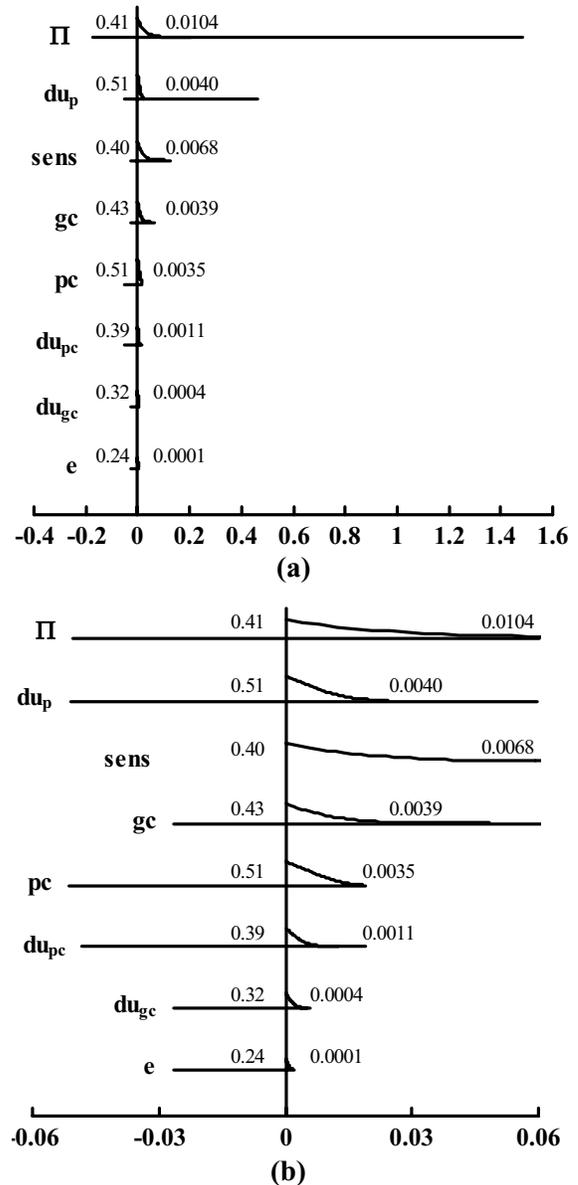


Figure 11 (a) A Javelin Diagram for the Giant Cell Arteritis Problem (b) Figure 11a Enlarged to Better Show the Javelin Heads



utility would increase from 0.837 to 0.842. Presumably this constitutes a significant increase, as otherwise it is unlikely Buchbinder and Detsky (1992) would have included this disutility in their model. Moreover, because this was included and no smaller values were considered, we may be justified in assuming that 0.005 is the smallest utility increment Buchbinder and Detsky (1992) considered significant. By this light, the

utility improvement 0.0104 for Π would be considered significant, as would the utility increment for *sens*, but no other individual parameter would possess significant information value. Note in particular that while variation in *gc* was likely to undermine the optimality of the BOA 43% of the time, its information value was only 0.0039. Where the tornado diagram of Figure 10 illustrates that the BOA is sensitive to variation in *gc*, the small information value suggests otherwise. Moreover, although the BOA appears vastly more sensitive to du_p than *sens* based on the potential payoff gains mapped out in Figure 10, the opposite is actually the case as $EUI_{sens} > EUI_{du_p}$ (0.0068 versus 0.0040), and the problem does not appear sensitive to du_p in information-value terms. This is typical of what we have found in our reexamination of many published sensitivity analyses (Felli and Hazen 1998, 1999), namely, that threshold proximity and probability of decision change frequently overestimate problem sensitivity compared to information value.

As in any probabilistic sensitivity analysis, the sensitivity conclusions obtained depend to some degree on the selected parameter distributions. However, in our experience, the shape of the distribution has only a small effect on probability or information-value conclusions. For example, in our reanalysis of 25 published decision analyses (Felli and Hazen 1999), there was little qualitative difference in probability or information-value conclusions when we replaced all parameter distributions by uniform distributions over the plausible parameter range. The changes in expected improvement consequent to using uniform rather than beta densities in Buchbinder and Detsky's (1992) problem are provided in Table 6.

Although there is deviation in information values depending on the choice of parameter density, the optimality of the BOA in Buchbinder and Detsky's (1992) problem appears sensitive only to variation in

sens and Π (based on the assumed demarcation point of 0.005 discussed earlier). The sensitivity conclusions remain the same regardless of whether parameter values are distributed according to the beta densities noted in Table 4 or uniformly distributed over their plausible ranges.

6. Summary

We have introduced javelin diagrams as a graphical aid for displaying results from a probabilistic sensitivity analysis. Like tornado diagrams, javelin diagrams provide the analyst with a compact display of payoff range information as a parameter or set of parameters are allowed to vary. In addition, javelin diagrams also provide the analyst with probabilistic and information-value-based indicators of decision sensitivity.

For ease of exposition, we assumed parametric independence in our examples. As it happens, correlations between parameters are rarely accounted for in conventional sensitivity analyses, even when they are well understood. For example, sensitivity and specificity of diagnostic tests are often treated as independent parameters for sensitivity analysis purposes despite their correlation through the thresholds used to declare positive test outcomes (Littenberg and Moses 1993). Some analysts build correlation directly into the model structure by specifying functional relationships between parameters. The Buchbinder and Detsky (1992) example we cite in §5 illustrates one such method common in the medical literature: employment of an efficacy parameter when drug and/or treatment choice can affect the likelihood of specific, downstream events. Dependencies built into model structure are accounted for by all conventional sensitivity analysis tools, as well as javelin diagrams. This may well be the best way to account for

Table 6 Information Values Are Noted for Parameters in the Buchbinder and Detsky (1992) Problem Using Two Different Distribution Assumptions

Density type	Parameters							Π
	<i>e</i>	<i>sens</i>	<i>gc</i>	<i>pc</i>	du_{gc}	du_p	du_{pc}	
Beta	0.0001	0.0068	0.0039	0.0035	0.0004	0.0040	0.0011	0.0104
Uniform	0.0003	0.0096	0.0038	0.0021	0.0001	0.0021	0.0003	0.0267

Notes. The beta values were obtained using the $\text{beta}(\alpha, \beta)$ densities noted in Table 4. The uniform values resulted from using uniform (*A, B*) densities, where *A* and *B* were the minimum and maximum parameter values listed in Table 4. Bold-faced values exceed 0.005.

probabilistic dependencies, as more explicit information on the joint densities of parameters is rarely available.

As is well known (e.g., Howard 1983), value-of-information computations can serve to guide the analyst’s model refinement and information acquisition choices in the decision-analysis modeling cycle. There is no reason why a javelin diagram could not be used at these earlier stages as well as the post hoc robustness stage we discuss in this paper. A small qualitative difference would occur for discrete variables, for which the javelin heads appear as step functions, but otherwise the same formulas and interpretations apply.

Appendix: Computing for Javelin Diagrams and Information Value

Let $a_0(\xi)$ be the optimal alternative as a function of ξ when all other parameters $\Pi \setminus \xi$ are held fixed at their base values $\Pi_0 \setminus \xi_0$. Computation of the expected improvement quantity

$$EVPI_\xi^0 = E_\xi[E[\Delta V_{a_0(\xi)} | \xi, \Pi_0 \setminus \xi_0]]$$

and the javelin curve

$$G_\xi(\Delta v) = P(E[\Delta V_{a_0(\xi)} | \xi, \Pi_0 \setminus \xi_0] > \Delta v), \quad \Delta v \geq 0,$$

is a straightforward exercise in Monte Carlo simulation, because the quantity $E[\Delta V_{a_0(\xi)} | \xi, \Pi_0 \setminus \xi_0]$ as a function of the random variable ξ is readily available from the decision problem formulation. The simulation can handle arbitrary parameter subsets $\xi \subset \Pi$, and if desired, Monte Carlo approximations for Π and multiple subsets ξ of Π can be calculated in a single simulation run.

As is well known, the mean of the nonnegative random variable $E[\Delta V_{a_0(\xi)} | \xi, \Pi_0 \setminus \xi_0]$ is the area under its complementary cumulative distribution $G_\xi(\Delta v)$. Therefore, the area under the javelin curve is $EVPI_\xi^0$.

If the regions $\{\xi | a_0(\xi) = a\}$ for each alternative a can be expressed as simple inequalities in ξ (as is often possible for single parameters ξ), then it may be possible to obtain closed-form expressions for $EVPI_\xi^0$ and the javelin curve $G_\xi(\Delta v)$. We have

$$EVPI_\xi^0 = E_\xi[E[\Delta V_{a_0(\xi)} | \xi, \Pi_0 \setminus \xi_0]] = \sum_a \int_{\xi: a_0(\xi)=a} E[\Delta V_a | \xi, \Pi_0 \setminus \xi_0] dF(\xi),$$

$$G_\xi(\Delta v) = P(E[\Delta V_{a_0(\xi)} | \xi, \Pi_0 \setminus \xi_0] > \Delta v)$$

$$= \sum_a P(E[\Delta V_a | \xi, \Pi_0 \setminus \xi_0] > \Delta v, a_0(\xi) = a).$$

For single parameters ξ and familiar distributions $F(\xi)$, terms in these expressions may often be evaluated in closed form.

For many decision problems, the expected improvement $EVPI_\xi^0$ is equal to the information value $EVPI_\xi$ for the parameter ξ . Here,

$$EVPI_\xi = E_\xi[E[\Delta V_{a^*(\xi)} | \xi]],$$

where $a^*(\xi)$ is the alternative a maximizing $E[V_a | \xi]$; $\Delta V_a = V_a - V_{a^*}$; and $a^* = a^*(\emptyset)$ is the overall optimal solution. Suppose, for example, that the function $E[\Delta V_a | \Pi]$ is multilinear in individual parameters $\xi \in \Pi$, that base values ξ_0 are equal to the means $E[\xi] = \bar{\xi}$, and parameters are probabilistically independent. Then, because $a^*(\xi)$ maximizes $E[V_a | \xi]$ and

$$E[V_a | \xi] = E_{\Pi \setminus \xi}[E[V_a | \xi, \Pi \setminus \xi] | \xi] = E[V_a | \xi, \bar{\Pi} \setminus \bar{\xi}] = E[V_a | \xi, \Pi_0 \setminus \xi_0],$$

it follows that $a^*(\xi) = a_0(\xi)$. Here we have used independence and multilinearity. Again, invoking independence and multilinearity, we obtain

$$EVPI_\xi = E_\xi[E[\Delta V_{a^*(\xi)} | \xi]] = E_\xi[E[\Delta V_{a_0(\xi)} | \xi]]$$

$$= E_\xi[E_{\Pi \setminus \xi}[E[\Delta V_{a_0(\xi)} | \xi, \Pi \setminus \xi] | \xi]]$$

$$= E_\xi[E[\Delta V_{a_0(\xi)} | \xi, \bar{\Pi} \setminus \bar{\xi}]]$$

$$= E_\xi[E[\Delta V_{a_0(\xi)} | \xi, \Pi_0 \setminus \xi_0]] = EVPI_\xi^0$$

as claimed.

References

Buchbinder, R., A. S. Detsky. 1992. Management of suspected giant cell arteritis: A decision analysis. *J. Rheumatology* **19**(8) 1220–1228.

Cawfield, J. D., M. C. Wu. 1993. Probabilistic sensitivity analysis for one-dimensional reactive transport in porous-media. *Water Resources Res.* **29**(3) 661–672.

Clemen, R. T. 1996. *Making Hard Decisions: An Introduction to Decision Analysis*. Duxbury Press, Belmont, CA.

Critchfield, G. G., K. E. Willard. 1986. Probabilistic analysis of decision trees using Monte Carlo simulation. *Medical Decision Making* **6**(2) 85–92.

Critchfield, G. G., Willard, K. E., D. P. Connelly. 1986. Probabilistic analysis methods for general decision models. *Comput. Biomedical Res.* **19** 254–265.

Doubilet, P., C. B. Begg, M. C. Weinstein, P. Braun, B. J. McNeil. 1985. Probabilistic sensitivity analysis using Monte Carlo simulation. *Medical Decision Making* **5**(2) 157–177.

Felli, J. C., G. B. Hazen. 1998. Sensitivity analysis and the expected value of perfect information. *Medical Decision Making* **18**(1) 95–109.

Felli, J. C., G. B. Hazen. 1999. Do sensitivity analyses really capture problem sensitivity? An empirical analysis based on information value. *Risk, Decision Policy* **4**(2) 79–98.

French, S. 1986. *Decision Theory: An Introduction to the Mathematics of Rationality*. Ellis Horwood, Chichester, U.K.

Goodman, C. A., P. G. Coleman, A. J. Mills. 1999. Cost-effectiveness of malaria control in sub-Saharan Africa. *Lancet* **354**(9176) 378–385.

Hamed, M. M. 1999. Probabilistic sensitivity analysis of public health risk assessment from contaminated soil. *J. Soil Contamination* **8**(3) 285–306.

Hazen, G. B., J. Sounderpandian. 1999. Lottery acquisition versus information acquisition: Prices and preference reversals. *J. Risk Uncertainty* **18**(2) 125–136.

- Howard, R. A. 1983. The evolution of decision analysis. R. A. Howard, J. E. Matheson, eds. 1989. *The Principles and Applications of Decision Analysis*. Strategic Decisions Group, Menlo Park, CA.
- Howard, R. A. 1988. Decision analysis: Practice and promise. *Management Sci.* **14**(6) 679–695.
- Keeney, R. L. 1992. *Value-Focused Thinking*. Harvard University Press, Cambridge, MA.
- Keeney, R. L., H. Raiffa. 1976. *Decisions with Multiple Objectives: Preferences and Value Trade-Offs*. 1993. John Wiley and Sons, Cambridge University Press, Cambridge, U.K.
- Lin, P. S., C. W. Chang, C. H. Juang. 1999. Probabilistic sensitivity analysis for one-dimensional solute transport. *Civil Engrg. Environmental Systems* **17**(1) 39–61.
- Littenberg, B., L. E. Moses. 1993. Estimating diagnostic accuracy from multiple conflicting reports: A new meta-analytic method. *Medical Decision Making* **13**(4) 313–321.
- Lord, J., M. A. Asante. 1999. Estimating uncertainty ranges for costs by the bootstrap procedure combined with probabilistic sensitivity analysis. *Health Econom.* **8**(4) 323–333.
- Manning, W. G., D. G. Fryback, M. C. Weinstein. 1996. Reflecting uncertainty in cost-effectiveness analysis. M. R. Gold, J. E. Siegel, L. B. Russell, M. C. Weinstein, eds. *Cost-Effectiveness in Health and Medicine*. Oxford University Press, Oxford, U.K.
- Murakami, Y., Y. Ohashi. 2001. Projected number of diabetic renal disease patients among insulin-dependent diabetes mellitus children in Japan using a Markov model with probabilistic sensitivity analysis. *Internat. J. Epidemiology* **30**(5) 1078–1083.
- Ng, A. K., J. C. Weeks, P. M. Mauch, K. M. Kuntz. 1999. Decision analysis on alternative treatment strategies for favorable-prognosis, early-stage Hodgkin's disease. *J. Clinical Oncology* **17**(11) 3577–3585.
- Pasta, D. J., J. L. Taylor, J. M. Henning. 1999. Probabilistic sensitivity analysis incorporating the bootstrap: An example comparing treatments for the eradication of *Helicobacter pylori*. *Medical Decision Making* **19**(3) 353–363.
- Piggott, J. H., J. D. Cawlfeld. 1996. Probabilistic sensitivity analysis for one-dimensional contaminant transport in the vadose zone. *J. Contaminant Hydrology* **24**(2) 97–115.
- Sisk, J. E., A. J. Moskowitz, W. Whang, J. D. Lin, D. S. Fedson, A. M. McBean, J. F. Plouffe, M. S. Cetron, J. C. Butler. 1997. Cost-effectiveness of vaccination against pneumococcal bacteremia among elderly people. *J. Amer. Medical Assoc.* **278**(16) 1333–1339.
- von Winterfeldt, D., W. Edwards. 1986. *Decision Analysis and Behavioral Research*. Cambridge University Press, Cambridge, U.K.
- Watson, S. R., D. M. Bued. 1987. *Decision Synthesis: The Principles and Practice of Decision Analysis*. Cambridge University Press, Cambridge, U.K.
- Williams, P., A. J. Dowson, A. M. Rapoport, J. Sawyer. 2001. The cost effectiveness of stratified care in the management of migraine. *Pharmacoeconomics* **19**(8) 819–829.

Copyright of Decision Analysis is the property of INFORMS: Institute for Operations Research and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.