

Logical Decisions^(R) for WindowsTM
Description
and
Demo Disk Example

What Logical Decision Does

Logical Decisions (LDW) is a sophisticated system for evaluating alternatives that differ on a number of evaluation variables or criteria. Alternatives can be anything you need to choose between -- jobs, potential employees, factory locations, or even what wine to have for dinner. LDW works best for decisions where you must consider many concerns at once, and where professional and value judgements will play a crucial role.

LDW lets you describe your alternatives with quantitative or qualitative **measures**. Measures are numerical or text variables that capture some quality of the alternatives. The measures can be anything that helps you distinguish between the alternatives -- salary, years of experience, distance from distribution centers, and alcohol content are all possible measures for describing alternatives. Each alternative has a raw score (called a **level**) for each measure.

LDW uses your judgements and preferences to rank the alternatives. LDW helps you review your preferences concerning the measures by guiding you through a series of questions. On the basis of your answers, Logical Decisions constructs the formula that ranks the alternatives.

In summary, you use LDW to complete the following steps:

- Define a set of alternatives to be ranked.
- Define measures to describe the alternatives.
- Enter the level on each measure for each alternative.
- Review your preferences so the measure levels can be combined.
- Rank the alternatives and choose the best one.

LDW uses powerful methods from the field of **Decision Analysis** to help you evaluate your alternatives. Decision analysis was developed in the 1960s and 1970s at Stanford, MIT and other major universities. In particular, Logical Decisions is based on the principles of **Multi-Attribute Utility Theory**, as described in "Decisions With Multiple Objectives: Preferences and Value Tradeoffs" by R. L. Keeney and H. Raiffa (John Wiley, 1976) and many other books. LDW's basic concepts are discussed briefly below. Then you can use the demo version to take a tour through LDW's many powerful features

Defining Alternatives.

While LDW gives you a lot of flexibility in the order you do things, you generally start an LDW analysis by defining your alternatives. You define an alternative by its name and its levels on the measures in a spreadsheet-like data matrix. If you already have a database of alternatives, LDW makes it easy to import them from several different formats.

Defining Measures

Measures are the variables that LDW uses to describe and then rank the alternatives. They are characteristics or attributes of the alternatives that have been made specific and that vary enough between the alternatives to influence the alternatives' desirability. You define measures in a spreadsheet-like data matrix or in a hierarchy of goals and measures like an organization chart.

After you have defined a measure, you enter its level for each of your alternatives. The level is the alternative's raw score or value for the measure. *Each* alternative has a level on *each* measure. The levels can be either numbers or brief text descriptions (such as "High", "Medium", or "Low").

Entering the Levels of the Measures

You enter measure levels in the spreadsheet-like Matrix window. You just type in a number or select from a list of possibilities for measures that use text labels.

You can also define measure levels in terms of probabilities. You can define the level of a measure for an alternative with any of six probability distributions.

Goals

LDW lets you group related measures under what is called a **goal**. Goals let you organize your concerns about your decision in a way that relates your broad overall concerns (hire the best employee) to more specific intermediate goals (hire the employee with the most experience) to specific measures (years of job experience and years of education). LDW lets you group goals and measures together under more general goals, forming a **goals hierarchy that looks like an organization chart**. You can rank your alternatives based on any of the goals or measures you have defined.

Reviewing Your Preferences (Weights)

In order to combine the levels of the measures into an overall score for an alternative, LDW needs to know the relative importances of different levels of the measures. Questions about relative importances (weights) require judgements that don't usually have a single correct answer. LDW helps you systematically review your preferences and values to identify a set of weights that is appropriate for you and for the particular decision you are analyzing.

Reviewing preferences is a three step process. First, you define a **preference set** to store your preference information. Then you (optionally) define how to convert measure levels to common units, and finally you define the relative importances of the measures and their interactions. LDW lets you define several Preference Sets for the same decision, so that you can see the effect of different points of view on the ranking of the alternatives. The next several paragraphs describe some of LDW's powerful methods for capturing preferences.

Converting the Measures to Common Units

LDW converts the levels on the measures to common units before combining them. The common units LDW uses are called **utility**. LDW assigns each measure's most preferred level a utility of 1 (one) and assigns each measure's least preferred level a utility of 0 (zero). LDW assigns intermediate levels using what is called a utility function, or more specifically a **SUF** (for **Single-measure Utility Function**). The shape of the utility function represents a set of preference judgements about the relative desirability of different levels on the measure.

The easiest way to convert levels to common units is with a straight line, where each unit change in a measure results in an equal change in utility. This is LDW's default. So, if you want a straight line conversion from a measure's nominal units to utility you don't have to do anything.

Sometimes a straight line conversion isn't right. Suppose you are hiring employees and "years of experience" is one of your measures. You might feel that after a certain point additional years of experience aren't that helpful. In that case you can use LDW's Assess::Common Units option to define a non-linear conversion.

Suppose 2 years experience is the least preferred level and 30 years is the most preferred level, but that you really want someone with around 5-10 years of experience. You might end up defining a curve that would make a change from 2 years experience to 3 years much more important to the ranking than a change of 29 to 30 years. LDW has options that let you easily define a (SUF) curve to describe these preferences.

LDW provides several options besides the SUF method for converting to common units. Several of these are based on the "Analytic Hierarchy Process". These options let you define the alternatives' utilities by making direct comparisons between them.

Defining the Relative Importances of Measures

LDW describes each alternative's overall desirability by a number called its **overall utility**. LDW computes an alternative's overall utility by combining each individual measure's utility using a weighted average. LDW can use a more complicated formula if there are interactions between the measures. Each measure's weight determines how much an alternative's utility on that measure will influence its overall utility.

LDW provides several different methods for assessing weights. These range from simple and easy to use to the most sophisticated methods available today.

The simplest method has LDW compute weights based on your ordering of the measures' importances (the "Smarter" method). Several other methods have you enter importance *ratios* for the measures (the "Smart" method and the "Analytic Hierarchy Process"). Finally, in the "tradeoff" method, LDW computes weights indirectly based on how much of one measure you would "trade off" to get more of another.

More on Tradeoffs

In LDW, a pair of simplified, alternatives used to help define weights is called a **tradeoff**. The alternatives highlight a trading, or give and take, between two measures. The first alternative has a preferred level on one measure and a less preferred level on the other, while the second alternative has a less preferred level on the first measure and a more preferred level on the second. Thus, if you traded the first alternative for the second, you would lose utility on the first measure and gain utility on the second. If the two alternatives are equally preferred then the decrease in the first measure just compensates for the increase in the second.

Here's an example. Suppose you are evaluating job candidates by years of education and years of experience. You might feel that a candidate with 14 years of education and 10 years experience is equally preferred to one with 16 years education and 3 years experience. Here, an increase of 2 years in education exactly compensates for a decrease of 7 years in experience. LDW uses this information to establish relative weights that will make the two candidates have equal overall utilities (assuming all the other measures are equal).

To define a weight for each measure, you define a *set* of tradeoffs, making sure that each measure appears in at least one tradeoff. LDW makes it easy to do this. LDW also provides many features that help you assess and review tradeoffs.

Interactions between measures

LDW can also model interactions between measures. Suppose you felt that education could compensate for experience and vice versa. In this case you might want a candidate with either 10 years of experience **or** 16 years of education to receive the highest utility on the combined measures. This is called **constructive interaction** in LDW. The methods for defining interactions are somewhat advanced and are discussed in the User's Manual.

Displaying Results

LDW provides you with many ways to review the results of your analysis. You can rank the alternatives based on overall utility or on of any other goal or measure. LDW sorts the alternatives by utility and displays them. The utilities generally range from a perfect utility of 1.0 to a lowest utility of 0.0

Besides ranking the alternatives, you can generate bar graphs showing the utilities for different alternatives, generate detailed comparisons between two alternatives, and analyze the effects of changes in the relative importances of the measures. You can

also create displays that show the effects of uncertainties on the ranking results. LDW uses Monte Carlo simulation to estimate the effects of uncertainties.

LDW also lets you display the formulas that it uses to generate the rankings.

Final Comments

LDW is a powerful tool to help analyze your decision problem, but it's only a tool. It isn't a black box that tells you what to do. It helps provide insights into the desirability of your alternatives and helps you think through difficult choices in a logical way. After you have ranked the alternatives and reviewed their differences, the reasons for the ranking results should be obvious. The results should reflect your preferences.

If you don't like the ranking, it often means that you could gain insights from more study of the results or that important considerations have not been captured in the measures.

The quality of the answers LDW generates depends critically on the quality of the data that it receives. In particular, you should take the time to

- Define your alternatives in detail.
- Make sure you have included all available alternatives.
- Think carefully about which measures are most appropriate for your alternatives.
- Make sure your measures cover all of the important considerations for choosing alternatives, that they are not redundant, that they can be measured for each alternative and that they are meaningful for decision making.
- Think carefully when defining SUF curves and answering the tradeoff questions.

A TOUR OF LOGICAL DECISIONS^(R) FOR WINDOWSTM

Introduction

This tour will use the Logical Decisions Demo Version to give you a feel for how Logical Decisions works and how it can give you insights into your most difficult decision problems. The Demo Version is a full working copy of LDW with the restrictions that you can't save anything to a file or print anything. The tour is divided into the following sections, that roughly parallel how you can use LDW to structure and solve a decision problem:

- Structuring the problem
- Assessing preferences
- Reviewing results
- Advanced Features

Step 1: Installing the LDW Demo Version

Insert the LDW Demo Installation disk 1 into your 3.5" drive. Go to the Run option in the Windows Program Manager or File Manager and type the command A:INSTALL. Follow the prompts and the LDW Demo Version will install itself into your Windows system. When you are done with the program you can select the Uninstall icon to remove the demo from your system.

Step 2: Loading a Logical Decisions file

We'll be working with a file called EXAMPLE.LDW. This file is an simplified evaluation of 90 MHZ Pentium computers and is based on a review that appeared in the Sept. 27 Issue of PC Magazine. To load in a logical decision file, start the Logical Decisions program and click on the leftmost button (or select the File::Open command). Select the file EXAMPLE.LDW from the list. LDW will open the file and display two windows initially -- a goals hierarchy and a spreadsheet matrix. We will take a look at these two windows next.

STRUCTURING A DECISION USING LDW

The first activity in analyzing a decision using LDW is to structure the problem. Here you will define the alternatives that you will be ranking, the variables (called "measures" in LDW) you will use to rank the alternatives and the goals you will use to organize the measures. In the file EXAMPLE.LDW, this step has been completed, but the tour will show you how this was done. Also see Step 28 for a look at how to import data from an outside source.

Step 3: Viewing a Summary of the Analysis

To get an overview of the Pentium computer analysis select the View::Summary option. You will see a dialog box that summarizes the analysis. You will see that we have 10 alternatives, 10 goals and 19 evaluation measures. This example has been simplified from the original, which contained 40 alternatives, 16 goals and 47 evaluation measures. You can see that LDW has the ability to handle quite large and sophisticated decision problems. There is no limit (other than memory) to the number of alternatives and measures you can have in an LDW analysis.

Step 4: The Matrix View

Select Matrix from the Window menu to display the Matrix view. If you have closed the matrix window, create a new one by selecting the View::Matrix option. The dialog box defaults are fine for now. The Matrix view is where you define the alternatives you will be ranking and describe the alternatives in terms of the measures. You will see a spreadsheet with alternatives along the left and measures along the top. Each cell in the matrix is the level (raw score) for an alternative on a measure. Note that, unlike many of its competitors, LDW uses natural units for each measure, such as dollars for price, MB for hard disk size, and the manufacturers names for Monitor Manufacturer.

Background: LDW is one of the few programs of its type that works with natural units (numbers and text descriptions). Other programs force a conversion to unnatural 1-10 scales or omit scales entirely, relying on direct comparisons of alternatives. These conversions obscure the objective part of the problem, the raw data. This makes it hard to gain consensus about the facts underlying a decision and can obscure the audit trail from the raw data to the final decision.

Step 5: Defining a Measure

Next we will look at what defines an evaluation measure. Each measure has a dialog box that you can see by double clicking on its cell at the top of its matrix column. The dialog box is where you describe each measure to LDW. Double-click on the "List Price" measure. The things to note here are the most and least preferred levels where you set the range for the measure and the cutoff level of \$4500. Computers that cost more than \$4500 will be marked as failures in several of the results displays.

Step 6: Defining Measures With Text Labels (NEW for 4.0)

Now lets look at the dialog box for a measure that uses labels. Double click on the "Case Style" measure cell. The important things in this dialog box are the "Use Labels" checkbox the "Define Labels" button. When you click on the "Define Labels" button you will see a dialog box with a list of the possible labels for the "Case Style" measure. These are ordered in decreasing desirability. This is clearly a subjective judgement, and you can use the buttons at the left of the dialog box to change the ordering or to add or delete labels. In step 18 we'll see how to assess the relative desirability of each label.

Background: Allowing text labels makes it easier to incorporate "soft" non-quantitative concerns into a decision analysis. It also makes it easier to import data from outside data files, that often use text fields to describe alternative

features and characteristics.

Step 7: Change a Number

Changing a number (called a "level" in LDW) is easy in the Matrix view. Just click on the desired cell and type in a new number. There is no requirement that the number you enter be within the range defined for the measure in its dialog box. In step 32 you can see how to define a measure level as a probability distribution. You can also import the measure levels as a group from an external database or spreadsheet file. You can see how to do this in step 31.

Step 8: Change a Label (NEW for 4.0)

Changing the level for a measure that uses text labels is also easy. If you click on a cell in the "Case Style" column, LDW will create a combo box in the cell for you to pick from. Just select the label you want, and LDW will make the change. Click on another cell to make the combo box go away.

Step 9: View the Goals Hierarchy

The second part of structuring a LDW decision is developing a "goals hierarchy". A goals hierarchy is a structure like an organization chart that links the evaluation measures to your overall goal, in our case "Buy the Best Computer". Goals are used to combine and link the measures. A goal is simply a placeholder where LDW aggregates the scores for the alternatives on the measures. The purpose of the goals hierarchy is to help ensure that you have captured all of the important concerns by evaluation measures, to simply preference assessments (as discussed in steps 19-23) and to let you rank on specific concerns. For example, we will be able to get a combined ranking for the alternatives on all the benchmark tests by ranking on the "Benchmarks" goal. To view the goals hierarchy, select Goals Hierarchy for Best Computer Goal from the Window menu. If you have closed this window, select the View::Goals Hierarchy option. The dialog box defaults will do for now. You may want to close the Matrix window as this will speed up the program somewhat. If the hierarchy is too big, you can use the Window::Zoom Out option (the seventh button from the left) to shrink it.

Background: LDW goals let you turn an ordinary data matrix into a sophisticated hierarchical database. For example, rather than worrying about seven different individual benchmarks or eight measures of case quality, you can compare the alternatives on the benchmarks or case goal to get a broad picture of how the alternatives perform in these areas. You can later "auger in" to get the details of why the alternatives rank as they do. The power of the goals is that they contain preference judgements. Different individuals might combine the variables under a goal in very different ways. They wouldn't be wrong, just different. Most databases have no way to incorporate preference data and thus can't aggregate their fields into meaningful higher level structures.

Step 10: Adding New Goals and Measures

The goals hierarchy is where you can add new goals to your analysis. To add a goal or a measure, first click on the goal you want to add under (say the "Company" goal) then press **Insert** or select the Edit::Insert option. LDW will ask if you want a new measure or a new goal. Select a measure, and LDW will create the new measure and show you its dialog box. Here you can change the measure's name and other parameters. When you click "OK," LDW will redraw the goals hierarchy and display the ellipse for the new measure. Select the Edit::Delete option after clicking on the new measure to delete it. You may need to reload the EXAMPLE.LDW file if you make too many changes by adding and deleting. You can also add measures and alternatives in the Matrix view using a similar process.

Step 11: Cutting and Pasting Goals and Measures

You can cut and paste goals and measures using the Edit::Cut, Edit::Copy, and Edit::Paste options. If you cut and paste a goal, all of the measures and goals beneath it are also cut and pasted. This makes it easy to quickly rearrange large parts of the goals hierarchy. If you try this, be sure to reload the EXAMPLE.LDW file before proceeding.

Step 12: Collapsing and Expanding the Goals Hierarchy

LDW has a unique feature that lets you expand and collapse the goals hierarchy like an outline. Just select a goal -- say

the "Company" goal -- (by clicking on it) and select the Hierarchy::Collapse option (the next to last button on the bar). LDW will redraw the hierarchy without showing the measures and goals below the "Company" goal. The line to the right of the goal indicates hidden goals and measures. To expand a goal, select the Hierarchy::Expand option (the last button). Note that this option only expands the hierarchy one more level.

ASSESSING PREFERENCES

Once you have completed the Matrix and Goals Hierarchy you are done with the structuring part of your analysis. You can come back and revise the structure later if you want to, but for now its time to move on. What you have now for the Pentium problem is a bushel of apples and oranges to compare. We have measures using a bewildering number of scales, ranging from the dollar scale of "List Price" to the descriptions of different case types to the unit-less scales used for the benchmarks. While this data is all fairly objective, there is no objectively correct way to combine it. We must make some subjective judgements about the relative importances of the different measures and goals. In fact, one of LDW's main purposes is to make a clear separation between the objective parts of the analysis (the structuring and measure levels) and the more subjective parts. Competing programs are generally much poorer at doing this because they don't let you use natural scales.

The purpose of the preference assessments is to guide you through the judgements needed to define the relative importances of the measures. We do this in two steps. First we assess how to convert the measure scales to common units (called utility) so we can combine them and then we assess the weights for the measures so we can give each measure and goal its proper importance. LDW provides more flexibility in doing this than any other program and promotes good assessments by helping you to remember the ranges of the measures, thus avoiding one of the most common preference assessment mistakes. Steps 13-20 will show you how the preference assessment process works.

Step 13: Preference Sets

LDW stores the information from the preference assessments in something called a preference set. Since different people or groups may have different preferences, LDW allows you to create many preference sets to reflect these differences. You can select a preference set as active with the View::Select a Preference Set option (the forth button from the left). All of the preference assessment options apply only to the active (selected) preference set. Each preference set has a dialog box that you can view with the Assess::Preference Sets option. The Pentium example has only one preference set, called "Pentium Buyer." The name of the active preference set is shown in the lower right of the LDW main window.

***Background:** Preference sets enforce the idea that we should be able to get consensus about the raw data for a problem (the numbers and text in the Matrix view spreadsheet), but that reasonable and even desirable to have different preferences and judgements about that data. Preference sets let you evaluate the same data set using different viewpoints and helps a group focus in on those judgements that make a critical difference in making a choice.*

Step 14: Show assessment Status (NEW for 4.0)

You can use the goals hierarchy window to help you with the preference assessment process. To do this, select the goals hierarchy window and then select the Hierarchy::Show Assessment Status option. This colors the hierarchy based on whether the preference assessment for each goal or measure has been completed. Measures or goals with completed assessments are drawn in red, measures or goals without completed assessments are drawn in red. Measures and goals not needing assessments are drawn in gray. Note that you can rank the alternatives at any time even if you haven't finished the preference assessments. LDW will use defaults (like equal weights) for any incomplete assessments.

Step 15: Assessing a "Utility Function"

The first step of the preference assessment is to convert the ranges for the individual measures to common units. LDW uses default common units called "Utility". Utility units (sometimes called "utils") range from 0.0 (worst) to (1.0) best. (you can use the Options::Names and Options::Utility Options to define a different name and range for the common

units). For continuous measures (like "List Price") LDW uses a continuous function to convert levels to common units. To see how this is done, select the "List price" measure in the goals hierarchy by clicking on it. Then select the Assess::Common Units option (the fifth button from the right). LDW will display the common units dialog box for "List Price". This dialog box controls how you will do the assessment for converting "List Price" to common units. The combo box labeled "Assessment Method" determines which method to use. We want the SUF (Single-measure Utility Function) method. Click "Assess" to begin the assessment. LDW will draw a graph showing how each "List Price" level (on the bottom axis) will be converted to utility (on the left axis). The green curve is the utility function. The purpose of the assessment is to determine the shape for the curve. If "a dollar is a dollar", then a straight line is appropriate. If we begin to get worried as the price gets too high we may want each extra dollar towards the right of the scale to cause a larger drop in utility than dollars at the left of the curve. You can do this informally by using our mouse to drag the green box at the center of the curve to the right. You can also change the curve by using the edit boxes at the bottom of the graph. Try entering 5250 into the "Level" edit box and pressing Enter. You can use a more formal assessment by choosing the SUF::Assess Value option (the third button from the right). In this option (NEW for 4.0), LDW asks you to identify a point that is halfway in terms of preference between the endpoints. That is, the change from each endpoint to the midpoint should be the same in terms of preference. The "In Depth" section of the manual describes how to do this assessment in more detail. When you are done, click "Done" and you will return to the SUF assessment window. Click "Done" again to return to the dialog box. The SUF menu provides other options for changing the shape of the SUF curve, including options to define a curve with many bends or with a midpoint more preferred than the endpoints.

Background: Most of LDW's competitors completely skip over this step. A straight line conversion to common units is assumed. Think of the problems this could cause with a variable like "Installed RAM", where increasing RAM from 4 mb to 8 mb can have a huge impact on performance but where increasing RAM from 32 to 36 mb (still a 4 mb increase) would probably have a negligible impact. LDW is the only software we know of that implements the "Assess Utility" and "Assess Value" options for assessing common units. These methods are widely taught in universities and MBA schools.

Step 16: Assessing Common Units with Text Labels (NEW for 4.0)

To see how the assessment process works for measures that use text labels, select the "Case Style" measure from the combo box at the top of the Common Units Dialog. The Assessment method should be "Direct Assessment." When you click "Assess," LDW will draw a bar graph showing the relative utility for each of the 4 possible labels for "Case Style". You can change the utilities by directly entering a number between 0 and 1.0 in the edit box for a label. You can also change the length of the bars by dragging at the right hand end. The cursor will change to a hand when you can drag. You have a lot of freedom in setting the utilities, but you should have one label get a utility of 1.0 and another get a utility of 0.0. Click "Done" to return to the Common Units Dialog box and "Done" again to return to the goals hierarchy. At this point the ellipses for all of the measures should be blue, indicating that we have completed the preference assessments for the measures.

Step 17: Assessing a tradeoff

Now we will turn our attention to assessing the relative importances (weights) for the measures. First we will assess a "tradeoff" to see how much we would be willing to pay for improved performance. To begin, Click on the "Best Computer" goal in the Goals Hierarchy and then select the Assess::Weights option (fourth button from the right). You will see the Assess Weights dialog box. The weight assessment stuff is on the left. You should make sure the "Best Computer" goal is selected in the combo box at the top. You can look at the "Assessment Method" combo box to see the different weight assessment methods available. Select the "tradeoff" method if it is not selected already. You also need to make sure that the "Use Representatives" box on the upper right is selected. Now click the left "Assess" button. You will see a dialog box that will help you select a pair of measures to compare in the tradeoff. You only have two choices, so they are already selected. Now click "Assess." LDW will now ask you to select a measure or goal to "represent" the "Bang" goal in the tradeoff. Select the "Winstone 94" measure from the (long) list. (Be patient, this should all make sense in a minute). Now LDW will draw a graph showing different levels of "List Price" and "Winstone 94". Our job is to identify two different pairs of levels (they'll be called A and B) that we prefer equally. LDW starts with an initial pair and asks us which we like best. A is expensive (\$6263) but has the best Winstone 94 performance (114.6). B is cheap

(\$2845) with poor Winstone 94 performance. We need to tell LDW which one we like best by clicking the proper button. Suppose we like B better. Click the B button and LDW let us improve A by making it cheaper. We want to improve A until the extra price for A just makes up for the improved performance. Suppose we are willing to pay \$500 for the better performance. Then the price for A should be \$3345 to make it equally preferred to A. Type 3345 into the edit box and press Enter. LDW will move A on the graph to reflect the new level. Now click "Equal" to say A and B are equally preferred. LDW can use this information to compute the relative weights for "List Price" and the "Bang" goal. Click "OK" and "Done" to return to the Assess Weights dialog box. To assess all the weights in the Pentium example you would need to assess 18 tradeoffs.

Background: LDW is the only software that implements the tradeoff method for preference assessments. It has the big advantage of forcing decision makers to consider the range of each variable as well as its name when answering weight assessment questions. Measures with a narrower range should receive less weight. Think of the how importance of list price compared to the other measures would change if its original range (\$2845 to \$6263) was reduced to a much narrower range (say, \$2845 to \$3845). The most common mistake in assigning weights to measures is to ignore the ranges of the measures and to assign weights based on the names of the measures alone. Tradeoffs also have the advantage of allowing the weights to be recomputed automatically if the ranges of the measures change.

Step 18: Assessing Weights With the Smarter Method

Now lets look at a simpler weight assessment method. Select the "Benchmarks" goal from the combo box at the top of the Assess Weights dialog box (or click on the benchmarks goal in the goals hierarchy and select the Assess::Weights option). Select the "Smarter" weight assessment method and click the left "Assess" button. You will see a list three of the benchmarks in the PC Magazine review, each with its range and an edit box. To complete the assessment we just need to rank order the measures in terms of importance. Suppose we think that the "Winstone 94" benchmark is most important. Put a 1 in its edit box. If we think the "Graphics Winmark" is next, give it a 2. Suppose we think that the "Disk Winmark" is least important. Give it a 3. Now click the "Done" button and LDW will compute some weights based on the our rank ordering. If you think the weights are OK, click "OK" to return to the Assess Weights dialog box.

Background: LDW is the only software that we know of that implements the "Smarter" method. It is probably the simplest weight assessment available.

Step 19: Assessing Weights With the Smart (Graphic) Method (NEW for 4.0)

Lets look at one more method. We'll use the Smart Method (Graphic) to assess the weights for the "Hard Drive" goal. First, select the "Hard Drive" goal in the combo box at the top of the Assess Weights dialog box. Next, select the "Smarter Method (graphic)" from the "Assessment Method" combo box and click the left "Assess" button. You will see a display similar to the Smarter assessment window, but with bars. The bars are proportional to the relative weights for the measures. You should think about which measure you would most like to change from its worst level to its best level. Suppose we would change the "Hard Disk Capacity" measure from 240MB to 2100 MB first. Assign this measure a "swing weight" of 100 by pulling its bar or by entering 100 in its edit box. Now think about which measure we would "swing" next. Suppose its the Disk Controller Loc(ation) measure. (If you have forgotten the labels for this measure, double click on its name to review its dialog box.) Suppose swinging it through its range provides 50% of the benefit as swinging the "Hard Disk Capacity" measure. Then assign "Disk Controller Loc(ation)" a swing weight of 50. Now click "Done" and LDW will adjust these swing weights so they sum to 1.0, as shown in the review screen. Click "OK" to return the Assess Weights dialog box.

Step 20: Interactions

If you've been wondering about the right side of the assess weights dialog box, that's where we assess "Interactions". Interactions are a unique and advanced LDW feature that let you use a formula other than a weighted average for combining the utilities for the measures. This is useful in situations where a good utility on one measure is enough for a good combined utility (maybe "Amount of Beer" and "Amount of Wine" for a party) or where a poor utility on one measure is enough to result in a poor overall utility ("Amount of Beer" and "Amount of Pretzels"). LDW provides 3 methods for

assessing interactions.

Background: LDW is the only software that allows interactions.

REVIEWING RESULTS

We have now reached the point where we have completed the problem structuring and the preference assessments. We now have all the information we need to rank the alternatives. LDW can not only rank the alternatives on any of the goals or measures (including the all important best computer goal) but can also provide a wealth of other displays that can provide considerable insights into why the alternatives rank the way they do. Steps 21-27 will show you a few of these displays. You can use these displays to gain insights to why the alternatives rank as they do. The goal of an LDW analysis is to eventually have our intuition match the numerical results. In retrospect, the results should seem obvious! If you don't agree with the results, it means you have not completely thought through the relative importances of the measures or that you have omitted an important concern or that you have not completely understood the ranking results. You will probably want to move back and forth between the results displays and the preference assessments (or even the structuring step) several times before you are comfortable with your analysis results. The results displays can then become very important for explaining how you reached your decision to others.

Step 21: Rank Alternatives

LDW's primary results display is the rank alternatives screen. Select the Results::Rank Alternatives option (the third button from the right). Select the default option to rank the alternatives by the "Best Computer" goal. The bar chart shows the overall ranking of the alternatives. The alternatives shown in white have failed at least one cutoff. You can find the details of the cutoff failures with the Results::Cutoff Summary option (NEW for 4.0).

Step 22: Modifying and Printing LDW Windows

You can modify and print all of LDW's Review and Results windows. You can click on text and graphics elements and drag them around. With options in the Edit menu you can modify text, change text fonts or enter new text. You can also cut, copy and paste graphic objects. With options in the Options menu you can change the color and patterns of text and graphic objects. The Options::Black & White option (fourth button on the right) lets you convert a picture to a patterned black and white image suitable for printing. The Window::Zoom In and Window::Zoom Out options (fifth and sixth buttons from the right (NEW for 4.0), let you quickly change the size of the picture. Many displays can be sorted with the Options::Sort (NEW for 4.0) option. You can copy an entire picture to the clipboard with the Edit::Copy option (of no individual graphics objects are selected). You can, of course, print any window, and you can save it to a WMF file with the File::Save Graphic option. Many table-like windows can be saved as ASCII text with the File::Save Text option.

Step 23: Stacked Bar Ranking (NEW for 4.0)

The Stacked bar ranking is a variation on the ranking results display. Here, each alternative's bar is made up of sub bars that represent the contribution to the utility from lower level measures and goals. Select the Results::Stacked Bar Ranking option. Select "Best Computer" from the list and select the "Sort Bars by -- Weight" and "Show Measures Under Goal" buttons. Enter 10 in the "Number of bars to stack" edit, then click "OK".

Step 24: Results Matrix

The results matrix is a spreadsheet-like display that shows the utilities for the alternatives on all the goals and measures. The first row of the spreadsheet shows the weights for the measures and goals. The spreadsheet is initially sorted by overall utility for the alternatives and by weight for the goals and measures. Select the Results::Ranking Results Matrix option to see this display.

Step 25: Compare Alternatives Graph (NEW for 4.0)

The compare alternatives graph lets you see the differences between any two alternatives in detail. Select the Results::Compare Alternatives Graph option and select the "Gateway 2000 PS 90" and "Zeos Pantera 90" alternatives to compare. In the graph, The measures favoring the Gateway are shown on the right in blue. The measures favoring the

Zeos are shown on the left in red. The bars are scaled in relation to the total difference, shown at the right in light blue. In this particular chart, the Gateway comes with a larger monitor, while the Zeos partially makes up for it with a lower list price. Use the [Results::Alternatives Comparison Table](#) option to see the numbers behind the graph.

Step 26: Scatter Diagram (Improved for 4.0)

The scatter diagram lets you graph the performance of the alternative on any two measures. For our example we can recreate PC Magazine's "bang for the buck" display by selecting the [Results::Sensitivity Graph](#) option. Select the "Bang" goal for the horizontal axis and the "List Price" measure for the vertical axis. enter 40 as the "Number of Alts to Show" and 10 as the "Number of Alts to Label." Keep 4 as the "Utility Lines to Show" and check the "Label Utility Lines" option. In the graph, the top 10 ranking computers are shown in color with their names. The other alternatives are shown as black crosses. The pairwise utility lines for "Bang" vs. "List Price" are shown in the background in light gray.

Step 27: Dynamic Sensitivity (NEW for 4.0)

The dynamic sensitivity options lets you quickly see the effect of changes in the weight for a measure or goal on the ranking of the alternatives. When you select the [Results::Dynamic Sensitivity](#) option, LDW creates two windows, one with the weights of the measures and goals and one with the overall ranking of the alternatives. You can temporarily change the weights by dragging on the ends of the weight bars. You can also enter weights directly by clicking on the weight number itself to get an edit box. When you are finished dragging or press enter, LDW will set the adjusted weight to its new value, proportionally adjust the other weights and re-rank the alternatives using the new weights.

OTHER LDW FEATURES

The next few steps demonstrate how LDW can import external datasets, how it works with multiple sets of preferences and how it works with probabilities and Monte-Carlo simulation

Step 28: Import applying cutoffs and limiting number (NEW for 4.0)

Next we will demonstrate how LDW can quickly screen a large database to quickly identify the most promising alternatives. Use the [File::Open](#) option to load the EXAMFRM.LDW file. Don't save the EXAMPLE.LDW file if LDW asks you. The EXAMFRM file is a file with the goals hierarchy and preference set from the EXAMPLE.LDW file and a single alternative, the "Ideal" pentium that scores best on all measures. This file could be created by using the "Import Structure" option under [File::Import](#) to identify the fields and measures from a database and then modifying that information to organize it into a goals hierarchy, establish cutoffs and to assess preferences. We can now use the structure in the EXAMFRM.LDW file to screen the alternatives in a tab delimited data file called PENTIUM.TAB. Select the [File::Import](#) option and select "Tab Delimited" from the list and the "Append to Current File" button. Next select the file PENTIUM.TAB from the file selection dialog box (it should be already selected). Next LDW will ask you to select the measures and measure categories to be updated. Just click "Select" to update everything. Say "No" to "Append records matching existing alternatives?," "Yes" to "Apply cutoffs?" and enter 5 as the "Number of alternatives to append." LDW will scan through the PENTIUM.TAB file, identify the 5 highest ranking alternatives that pass the cutoffs and append them to the current analysis. You will see this reflected in the updated matrix display.

Step 29: Preference Set Summary (NEW for 4.0)

Use the [File::Open](#) option to load the file TRUCKS.LDW. Say no if LDW asks you about saving files. Next select the [Results::Preference Set Summary](#) option. Take the default options in the dialog box. LDW will create a table showing the overall utility for each of the three preference sets in the file. LDW also shows the mean and standard deviation over the utilities for each alternative.

Then next three steps show how to define probabilities and show probabilistic results in LDW.

Step 30: Define a probability

Use the [Window::Matrix](#) option to select the Matrix view window. Next double click on the cell for "Resale Value" for the

"Mitsubishi" alternative. You will see LDW's define probability dialog box. This is where you define probability distributions over measure levels. You can see that the Mitsubishi has been assigned a normal (bell curve) distribution over resale value. Click the "Modify" button to see the mean and standard deviation parameters needed to define a normal distribution. Each of the five types of distribution available in LDW requires slightly different parameters. Next, click OK to return to the matrix. You can also define probabilities over labels. Double click on the "Styling" cell for the "Mitsubishi." As you click on the different labels in the list, you will see the probability assigned to each one.

Step 31: Showing probabilities in Rank Alternatives (NEW for 4.0)

You can see the influence of the probability distributions on the results in the [Results::Rank Alternatives](#) option. Select the option, select the "Best Truck" goal and also check the "Show Uncertainties" box. You will see a modified ranking results chart where each alternative has a second (black) bar that indicates the range of overall utilities that might occur, given the probabilistic levels defined in the matrix. LDW computes these ranges using Monte-Carlo simulation to sample from the probability distributions and compute a set of possible utilities based on those samples. You can control the simulation with the [Options::Simulation Options](#) option. LDW performs the simulation just before displaying the results.

Step 32: Probability graph (NEW for 4.0)

To see the uncertainty for a single alternative in detail, select the [Results::Alternative Uncertainty Graph](#) option. Select the "Mitsubishi" alternative and the "Best Truck" goal and use the defaults for the other options. LDW displays a graph showing the cumulative probability distribution over overall utility for the Mitsubishi, and also shows a histogram showing where the simulation results fell.

Step 33: Probability table (NEW for 4.0)

To compare the uncertainties for all the alternatives over a single measure or goal, use the [Results::Uncertainty Summary](#) option. Select the default of "Best Truck Goal" in the dialog box and LDW will display a table showing the uncertainties in overall utility for each of the alternatives, given the simulation results. The columns should be fairly self explanatory.

This concludes the tour of the Logical Decisions program. We hope you enjoyed it. There are still many options we haven't covered, so feel free to explore on your own. There are also several other LDW files included with the demo version. These include AHP.LDW -- a simple buying a house example done using the Analytic Hierarchy Process; HOUSEEX.LDW -- a much more sophisticated house example; RESEX.LDW -- an evaluation of alternative reservoir locations based on an actual project.

Logical Decisions is dedicated to providing state-of-the-art decision analysis tools for businesses, academic institutions, government agencies and individuals. Logical Decisions also provides decision analysis consulting and training. Call us to discuss your requirements.

Logical Decisions
1014 Wood Lily Drive, Golden, CO 80401
1-800-35-LOGIC Fax (303) 526-7537