# Multi-Criteria Decision Thinking

## Idea In Short

Before comparing options, agree the criteria, their weights, and any must-have thresholds in writing. Then score every option against each criterion, calculate a weighted total, and test whether a different but reasonable set of weights changes the winner. If it does, the real decision lies in the weights, and leadership should debate them openly. Screen out any option that fails a non-negotiable threshold before scoring begins, so that a strong result on one criterion cannot hide a fatal flaw on another. Use this method when a decision involves several objectives that pull in different directions and when stakeholders need to see how the choice was reached.

Few significant decisions involve a single objective. A company picking a site for a new distribution center cares about cost, delivery speed, labor supply, room to expand, and exposure to disruption, and those goals rarely favor the same location. Without a method, the loudest voice or the most recent presentation tends to win. Multi-Criteria Decision Analysis [MCDA] gives the group a structure for weighing competing goals in the open. This article explains how the method works, shows a worked comparison, and describes the places where scoring misleads.

## Why one number is rarely enough

A decision with one objective is easy. The traveler who wants the cheapest flight from New York to San Francisco finds the lowest fare and buys it. Most real choices add more: a convenient departure time, a direct route, an aisle seat, a carrier with a strong safety record, and frequent flyer miles.1 Organizations face the same tension at larger scale and with more at stake. A hospital choosing a scheduling system, a ministry allocating a grant, and a manufacturer choosing a supplier all pursue several aims that conflict.

When objectives conflict, intuition tends to simplify the problem by focusing on one or two aims and ignoring the rest. The result can feel decisive, but it hides the trade-offs that

others may want to examine. MCDA forces the trade-offs into view. It asks the team to name every objective, say how much each matters, and show how each option performs.

## The anatomy of the method

The British government publishes guidance for analysts who wish to apply MCDA in policy development under its Green Book, and it revised that note in May 2024, replacing earlier multi-criteria analysis guidance.2 The steps it describes are common to most versions of the method. Analysts define the decision and the options, set the criteria, measure each option's performance, assign weights, calculate weighted scores, and test the results. The method goes by several names, including multiple attribute utility theory and multi-objective decision analysis, and it is used in business, government, and medicine.3

The simplest form combines criteria with a weighted linear function. Each option's score on a criterion is multiplied by that criterion's weight, and the products are added. More elaborate versions allow for nonlinear value curves, but the simple form handles most business decisions well, and it has the advantage that any colleague can check the arithmetic.

## Choosing criteria

Good criteria are relevant, measurable, and independent of each other. Relevant means that each one ties to an objective the decision serves. Measurable means that someone can assign a score using evidence rather than mood, even when the scale is qualitative. Independent means that two criteria do not capture the same thing under different labels. A team that scores both delivery speed and customer wait time has counted one factor twice.

### Weights are value judgments

Weights express how much the organization cares about each objective, and no analysis can derive them from data. A sound approach asks decision makers to consider the range of each criterion across the options and judge which swing matters most. A difference of one day in delivery time may be worth more than a difference of ten percent in cost, or the reverse, depending on the strategy. Record who set the weights and why. When stakeholders disagree, run the analysis with each set of weights and see whether the ranking changes. Often it does not, which ends the argument. When it does, the team has

located the real disagreement.

## A worked comparison

Consider a company choosing among three sites for a distribution center. The team sets five weights: cost 35 percent, service speed 25 percent, labor availability 20 percent, room to expand 10 percent, and disruption risk 10 percent. Scores run from 0 to 100, with higher numbers better. Site A scores 80, 60, 50, 40, and 70 on those criteria. Site B scores 60, 80, 70, 70, and 60. Site C scores 90, 40, 60, 90, and 50.

The weighted totals are 64 for Site A, 68 for Site B, and 67.5 for Site C. Site B leads by half a point over Site C, a margin too thin to settle anything. Now test an alternative in which a cost-focused board raises the weight on cost to 45 percent and cuts service speed to 15 percent. Site A and Site B tie at 66, and Site C jumps to 72.5. The ranking depends on one weighting judgment, and the executives now know precisely what to debate.

## Even swaps as an alternative

Some teams distrust explicit weights because the numbers feel arbitrary. Hammond, Keeney, and Raiffa proposed a method called even swaps, which gives a practical way to make trade-offs among objectives across a range of alternatives.4 The idea is to adjust two options until they match on one criterion, compensating with a change on another criterion that the decision maker judges equivalent. Criteria on which all options now tie drop out, and so do options that another option beats on every remaining criterion. The authors are candid that the method will not make complex decisions easy, since the user must still decide what a given trade is worth. What it does is simplify the mechanics so that attention goes to those value judgments.

## Use in public and corporate decisions

Public bodies value MCDA because it leaves a record. The Scottish Government's prevention toolkit describes a prioritization framework such as MCDA as a robust method for comparing policy options, producing a consistently scored ranking and a transparent record of criteria, assumptions, and trade-offs.5 The same toolkit notes that the method can be used lightly, to show the factors in a decision quickly, or in extensive models. Corporate teams gain the same benefit. A documented analysis helps when a board, an auditor, or a

losing bidder asks why one option beat another, and it lets the next team start from tested criteria instead of a blank page.

## Where scoring misleads

The weighted-sum approach is compensatory, which means a strong score on one criterion can offset a weak score on another. That behavior suits many decisions and ruins some. A supplier that excels on price and delivery but fails a safety audit should not win on the arithmetic. Set thresholds first and exclude any option that fails them. Other hazards include scales that differ in spacing, so that ten points means different things on different criteria, and weights adjusted after the first results appear to favor a preferred option. Prevent the second by recording the weights before scoring begins and logging any change with a reason.

## A routine for teams

A repeatable sequence keeps the analysis honest and quick. For a mid-sized decision, the whole exercise can fit in a half-day workshop, and the worksheet it produces lets later reviewers see how the choice was made. The steps are:

1. Define the decision, the options, and the must-have thresholds
2. Choose five to eight independent criteria and define how each will be measured
3. Agree the weights before scoring and record who set them
4. Score each option using evidence and note the source
5. Calculate totals, then rerun with at least two alternative weight sets

Present the final recommendation along with the sensitivity results. If the winner changes under reasonable weights, say so plainly and ask the accountable executive to decide among the weight sets. That conversation is the real decision, and it deserves to be held in the open.

## Limits of the method

MCDA cannot repair a poorly framed decision, and it cannot make the unmeasurable measurable. Scores for qualitative criteria such as strategic fit remain judgments, however tidy the spreadsheet looks. Analysts should therefore resist the temptation to equate a

decimal place with insight, and should present totals as rounded figures alongside the evidence that supports each score. The method also adds effort that small, reversible decisions do not warrant, so a lightweight version with three criteria and a rough weighting is often sufficient for those. Finally, the exercise works only when the accountable executive is prepared to act on the result, because a team that uses the analysis to justify a conclusion already reached will quickly stop trusting it.

- 1Even swaps: a rational method for making trade-offs
- 2Use of multi-criteria decision analysis in options appraisal of economic cases
- 3Multiple-criteria decision analysis
- 4Even swaps: a rational method for making trade-offs
- 5Prevention toolkit

## Summary

Multi-criteria decision analysis does not remove judgment. It relocates judgment to the places where a team can examine it: which criteria count, how much each one counts, and how well each option performs. The weighted total is the least important output. The documented trade-offs, the sensitivity tests, and the shared understanding of why one option won are the lasting benefits. Teams that fix thresholds first, keep criteria independent, and publish their weights tend to reach decisions that survive scrutiny from boards, regulators, and the people affected. Pick a pending choice, write five criteria, assign weights, and see how much disagreement the exercise surfaces.