Sensitivity Analysis

Determining the appropriate values to assign to the numerical items (the **input variables**) in a model is a critical and challenging part of the model building process in decision analysis. But finding numerical values for real problems requires gathering relevant data, which can sometimes be difficult. So we often make do with rough estimates. Because of the uncertainty about the true value of a numerical item, it is important to find out how the solution derived from the model would change (if at all) if the numerical value assigned were changed to other plausible values.

This process is referred to as **sensitivity analysis**. We can apply such analysis to the various ingredients that go into the calculation of the payoffs, to the probabilities of the different states of nature and even to the decision maker's utility function.

In essence, sensitivity analysis answers the question, "What makes a difference in this decision?" The aim of modelling in decision making is to produce a **requisite decision model** - one whose form and content are just sufficient to solve a particular problem. That is, the issues that are addressed in a requisite decision model are the ones that matter, and those issues left out are the ones that do not matter.

Alternatives can be screened on the basis of deterministic and stochastic dominance, and inferior alternatives can be eliminated. Identifying dominant alternatives can be viewed as a version of sensitivity analysis, for use early in an analysis. In sensitivity analysis terms, analysing alternatives for dominance amounts to asking whether there is any way that one alternative could end up being better than a second. If not, then the first alternative is dominated by the second and can be ignored. So sensitivity analysis can lead to modifying the structure of a model.

One-Way Sensitivity Analysis

Consider the "Build Big - Build Small" example we looked at earlier in the notes. The sensitivity analysis question in this case is, what items really make a difference in terms of the decision at hand? For example, do the various capital costs really matter? If the annual income obtainable during the first two years from a small plant in a high demand market changes by some amount, will that impact our initial decision? We can begin to address questions like these with one-way sensitivity analysis.

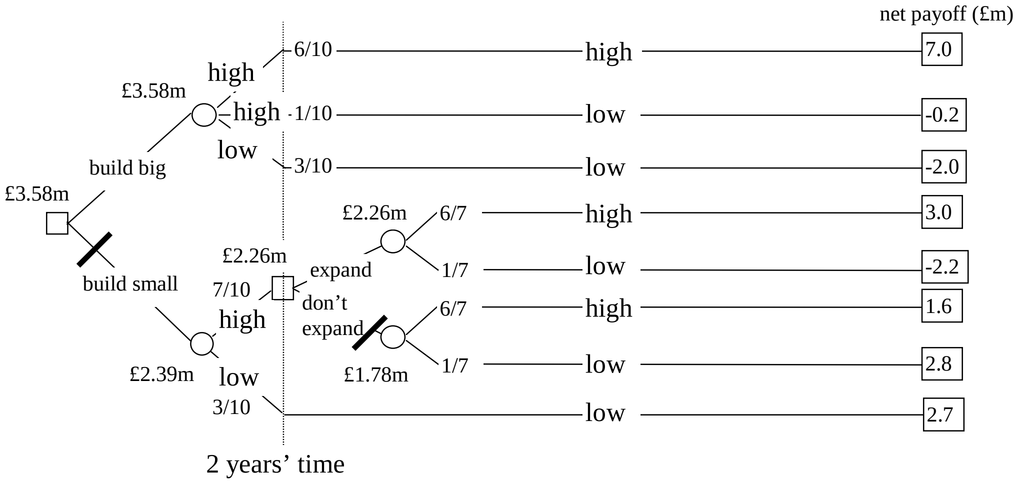

(Note that the expected payoff figures on the "build small" subtree differ slightly from those given earlier in the notes. This is due to rounding error. These ones are more accurate!)

The following table gives the capital cost and annual income figures (all given in £m), estimated at the base levels that were assumed in our previous analysis of the problem In addition, it includes reasonable upper and lower bounds that represent the decision maker's ideas about how high and how low each of these items might be. He might specify upper and lower bounds as absolute extremes, beyond which he is absolutely sure that the variable cannot fall. Another approach would be to specify the bounds such that he would be "very surprised" (a l-in-10 chance, say) that the variable would fall outside the bounds.

Input Variable	Base Value	Lower Bound	Upper Bound
Capital Costs:			
Building Big plant	3.00	2.00	4.50
Building Small Plant	1.30	1.00	1.50
Expanding Small Plant	2.20	1.50	3.00
Annual Income:			
Big Plant - High Demand - years 1 to 10	1.00	0.80	1.20
Big Plant - Low Demand - years 1 to 10	0.10	-0.10	0.25
Small Plant - Low Demand - years 1 to 10	0.40	0.20	0.60
Small Plant - High Demand - years 1 to 2	0.45	0.30	0.60
Small Plant - High Demand - years 3 to 10	0.25	0.10	0.40
Expanded Plant - High Demand - years 3 to 10	0.70	0.50	1.00
Expanded Plant - Low Demand - years 3 to 10	0.05	-0.10	0.15

Consider the capital cost of building a big plant. From the table, we see that the decision maker is not at all sure what the cost might turn out to be, and that it can vary from £2m to £4.5m. What does this imply for the initial decision? The simplest way to answer this question is with a **one-way sensitivity graph** as shown below.

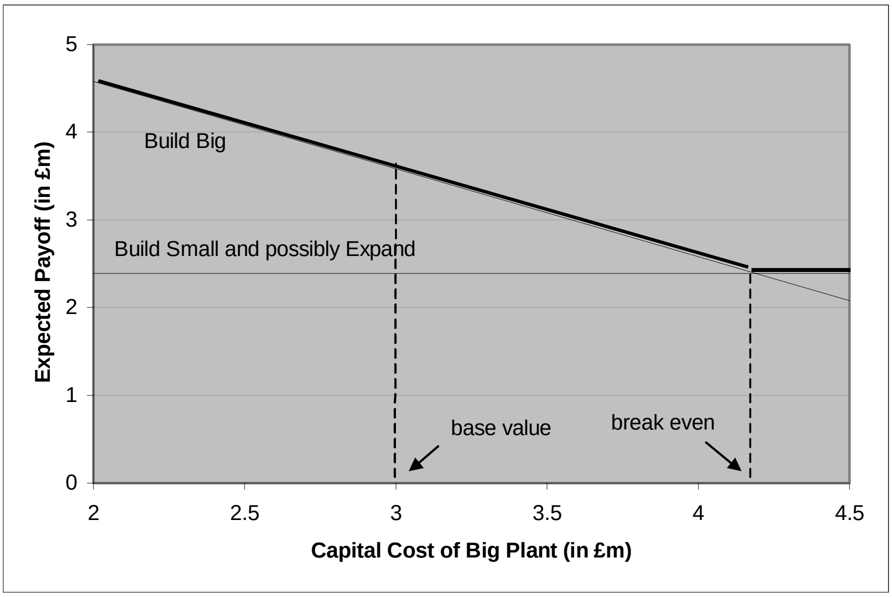

The downward-sloping line shows expected payoff from building big as the capital cost of building a big plant varies from £2m to £4.5. The horizontal line represents expected payoff from initially building small and then possibly expanding after two years. (Obviously the payoff from the latter action does **not** depend on the cost of a big plant.) The point where these lines cross is the threshold at which the two alternatives each yield the same expected payoff (£2.39m), which occurs when the cost of building big equals £4.19m.

The heavy line indicates the maximum expected payoff the decision maker could obtain at different values of capital cost of building a big plant, and the different segments of this line are associated with different strategies (Build Big versus Build Small and possibly expand).

Tornado Diagrams

The output value of interest and the one that determines our initial decision is actually neither the expected payoff from building big nor the expected payoff from building small and possibly expanding, but rather the difference between these two. We have the decision rule:

If *Expected Payoff from Building Big - Expected Payoff from Building Small > 0*

Then *Build Big*

Else *Build Small*

A **tornado diagram** allows us simultaneously to compare one-way sensitivity analysis for many input variables and a single output variable. Suppose we take each input variable in the above table and allow it to vary between its high and low values to determine how much change is induced in the difference between expected payoffs.

The following graph graphically shows how the difference between expected payoffs varies as the input variables are independently varied between the high and low values. For instance, ***with everything else held at the base value***, setting the annual income from a big plant in a market with high demand at £0.8m instead of £1m implies a difference between expected payoffs of -£0.05. This is plotted on the graph as the left end of the bar labelled income(Big plant - High demand). On the other hand, setting this income at the high end of its range, £1.2m, leads to a difference between expected payoffs of £2.43m. Thus, the right end of the bar is £2.43m.

We follow this same procedure for each input variable. The length of the bar for any given input variable represents the extent to which the difference between expected payoffs is sensitive to this variable. The graph is laid out so that the most sensitive variable - the one with the longest bar - is at the top, and the least sensitive is at the bottom. With the bars arranged in this order, it is easy to see why the graph is called a tornado diagram.

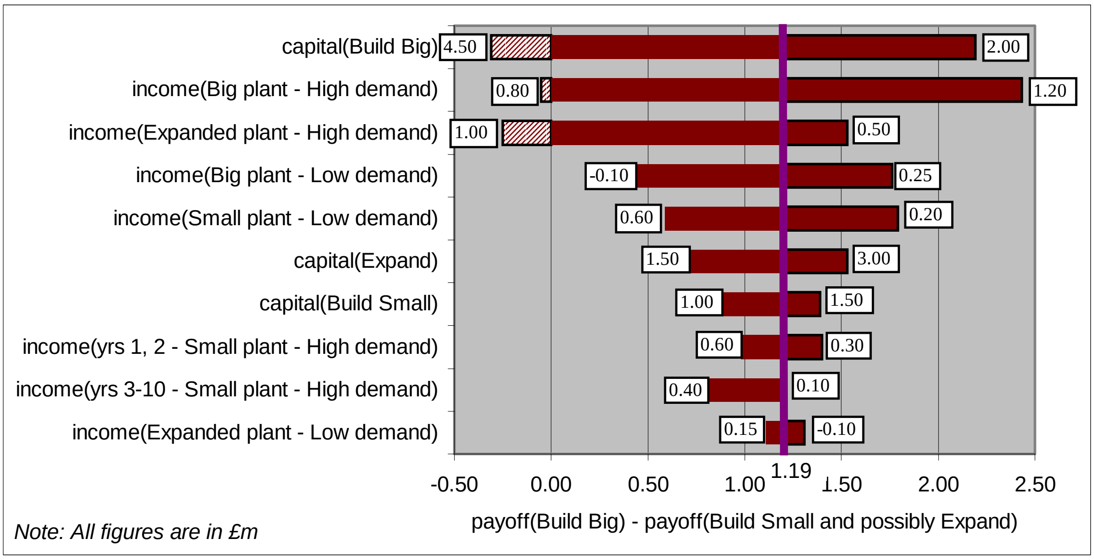

The vertical line at £1.19m represents the difference in expected payoff when all the input variables take their base value.

Interesting insights can be gleaned from this diagram. For example, the decision maker's uncertainties regarding the capital cost of building a big plant and the incomes of a big plant or an expanded plant in a high demand market are extremely important. They all have substantial effects on the difference in expected payoff, and the bars for these three variables cross the critical £0m line, below which building small and possibly expanding is the preferred strategy.

On the other hand, the difference in expected payoff is very insensitive to the capital cost of building a small plant as well as to the incomes of a small plant in a high demand market and an expanded plant in a low demand market.

The tornado diagram tells us which variables we need to consider more closely (perhaps making an effort to obtain more accurate estimates) and which ones we can leave at their base values. In further analysing this decision we simply can leave many of these input variables at their base values.

Two-Way Sensitivity Analysis

The tornado-diagram analysis provides considerable insights, although these are limited to what happens when only one input variable changes at a time. Suppose we wanted to explore the impact of several variables at one time? Although this can sometimes be tricky, it is easy to study the joint impact of changes in the two most critical input variables (the capital cost of building a big plant and the income of a big plant in a high demand market) since neither of these input variables has an impact on the expected payoff of building small and possibly expanding.

Referring to the two critical input variables as Cost and Income:

expected payoff from building big = 0.6 x 10 x Income + 0.1 x (2 x Income + 8 x 0.1) + 0.3 x 10 x 0.1 - Cost
= 6.2 x Income + 0.38 - Cost

expected payoff from building small and possibly expanding = 2.39

difference in expected payoffs = 6.2 x Income - Cost - 2.01

Imagine a rectangular space that represents all of the possible values that the two input variables Cost and Income can take. The straight line graph for

Cost = 6.2 x Income - 2.01

splits this space into three regions:

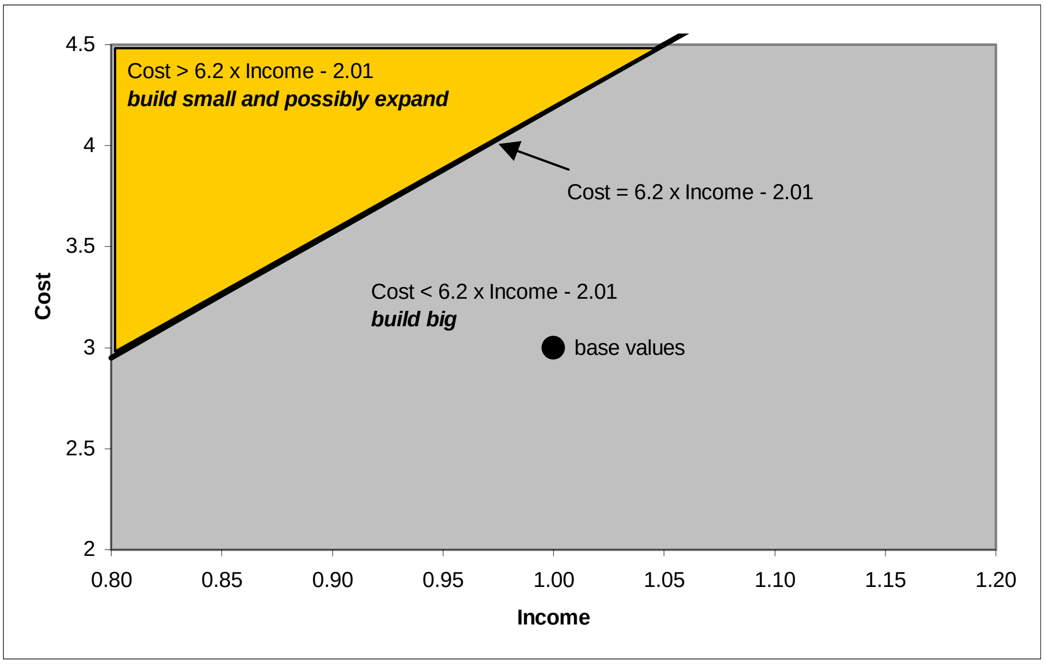

On the line, the difference between expected payoffs is zero and it does not matter which strategy is adopted. For points below the line, Cost < 6.2 x Income -2.01. Here, the expected payoff of building big is the greater and the decision maker should build big. Above the line, he should build small and possibly expand.

The point labelled “Base Values” shows that when we plug in the base values for the Income and Cost variables, we find that it is best to build big. It would take quite a large change, in both variables to get to a situation where the other strategy was preferable.

Obtaining an analytical solution may not be so easy with more complicated decision problems and we might then need to use a spreadsheet “two-way table” to carry out such a sensitivity analysis.

Sensitivity to Probabilities

The next step in our analysis will be to model the uncertainty surrounding the levels of demand. The base value risk figures were originally given in the form of joint probabilities:

Prob(Initially High, sustained High) = 6/10
Prob(Initially High, long-term Low) = 1/10
Prob(Initially Low, continuing Low) = 3/10
Prob(Initially Low, long-term High) = 0%

Prob(long-term Low) = 4/10

We can express the base value probability instead as:

prob(Initially High) = 7/10	prob(Sustained High | Initially High) = 6/7
prob(Initially Low) = 3/10	prob(Long-term Low | Initially High) = 1/7
	prob(Continuing Low | Initially Low) = 1
	prob(Long-term High | Initially Low) = 0

Let us assume that there is no doubt that if the demand is initially low it will remain so. (Otherwise we might need to change the structure of the solution to allow for the possibility of expanding in the face of an initial low demand.) Then there are actually only two distinct probability figures:

p = prob(Initially High)   q = prob(Sustained High | Initially High)

The other two can be calculated from p and q:

prob(Initially Low) = 1 - p   prob(Long-term Low | Initially High) = 1 - q

We can conduct one-way sensitivity analysis to investigate how expected payoff varies with either p or q, keeping the other one at its base value. For example, allowing p to vary but keeping q at its base value of 6/7, we get the decision tree:

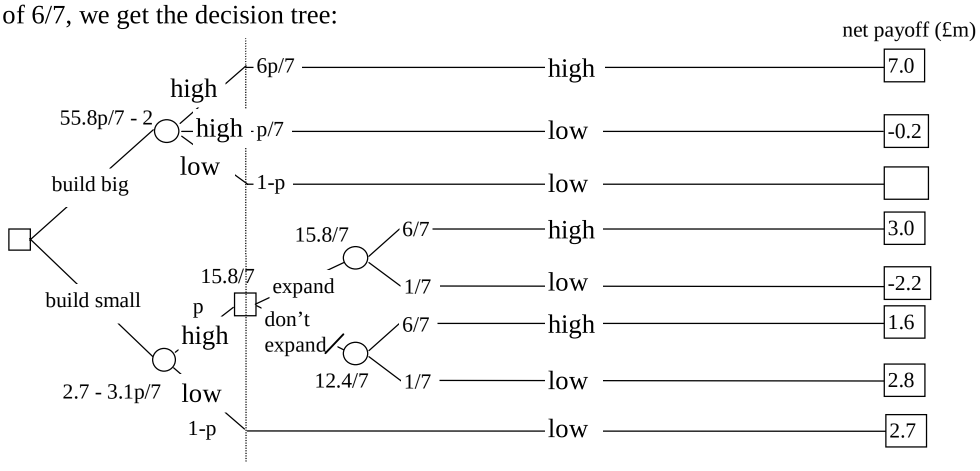

In the event of initial high demand, the expected payoff from expanding is higher than the expected payoff from not expanding, and so we can prune the latter action.

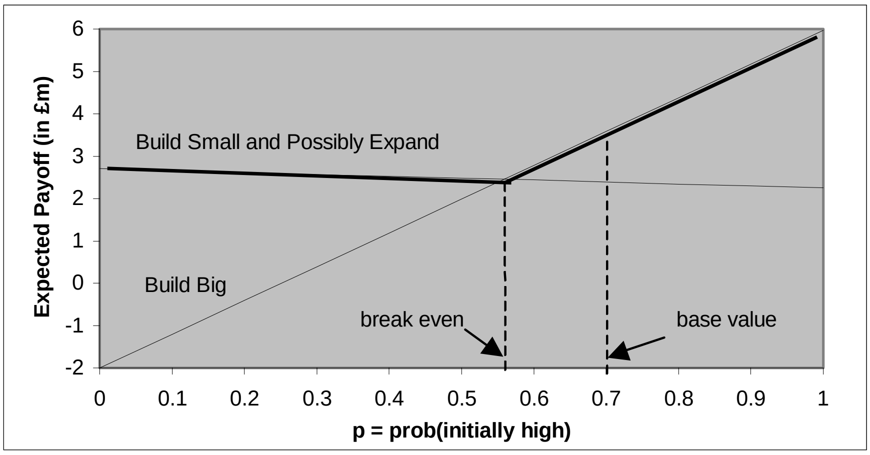

The break-even value of p can be calculated by solving:

$$55.8p/7 - 2 = 2.7 - 3.1p/7$$

i.e. $58.9p/7 = 4.7$

i.e. $p = 0.56$

Below this level it seems that its is best to play safe and build small, only expanding if the initial demand is high. Above this level it is best to build big.

Two-way sensitivity analysis involves investigating the expected payoff figures whilst both p and q are allowed to vary over their full ranges of 0 to 1. Here is the decision tree with both p and q as variables:

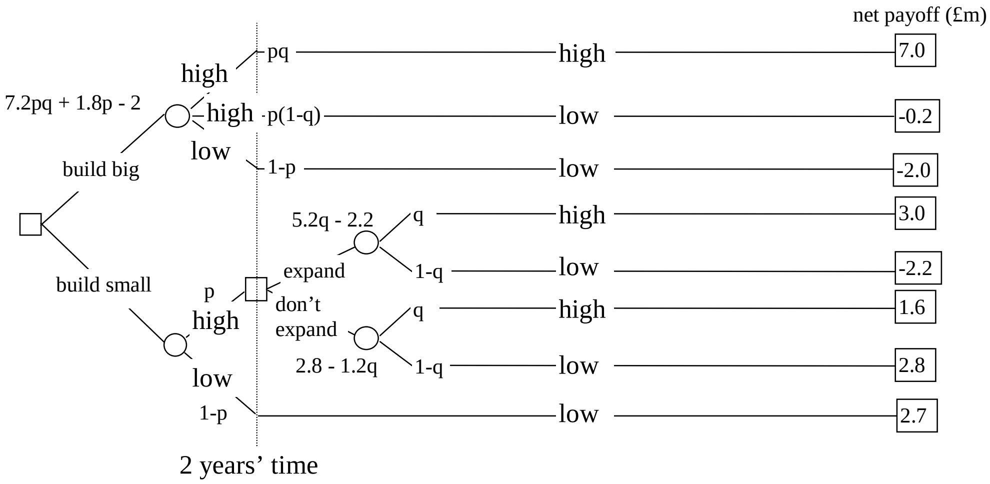

If we first compare the expected payoff from expanding or not expanding, in the light of initial high demand, we can calculate the break-even value of q from:

$$5.2q - 2.2 = 2.8 - 1.2q$$

i.e. $6.4q = 5.0$

i.e. $q = 0.78$

<u>If q > 0.78</u>: it is best to expand if the initial demand is high, and the expected payoff from building small is:

$$p \times (5.2q - 2.2) + (1 - p) \times 2.7 = 5.2pq - 4.9p + 2.7$$

Break-even between building big and building small is then given by:

$$7.2pq + 1.8p - 2 = 5.2pq - 4.9p + 2.7$$

i.e. $2pq + 6.7p = 4.7$

i.e. $\mathbf{p = 4.7/(2q + 6.7)}$

<u>If q < 0.78</u>: it is best not to expand even if the initial demand is high, and the expected payoff from building small is:

$$p \times (2.8 - 1.2q) + (1 - p) \times 2.7 = 0.1p - 1.2pq + 2.7$$

Break-even between building big and building small is then given by:

$$7.2pq + 1.8p - 2 = 0.1p - 1.2pq + 2.7$$

i.e. $8.4pq + 1.7p = 4.7$

i.e. $\mathbf{p = 4.7/(8.4q +1.7)}$

We can plot the two curves of p in terms of q, in the region over which p and q vary between 0 and 1.

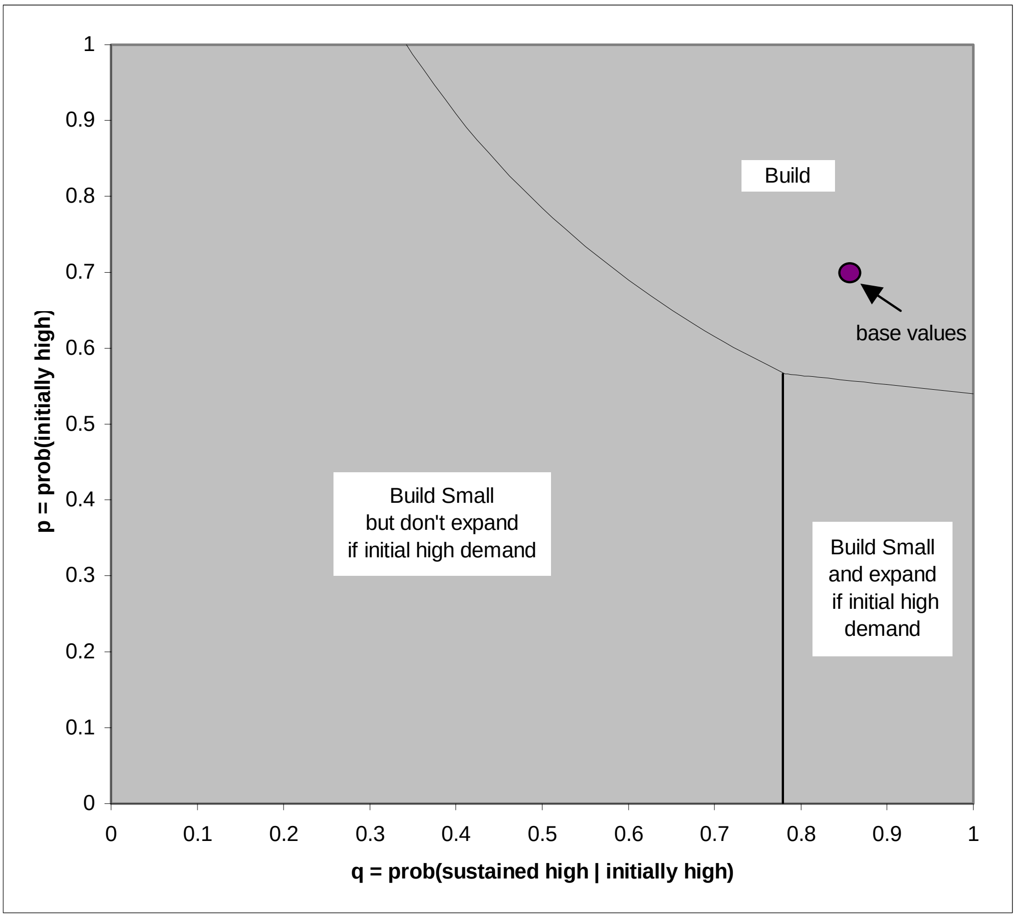

Points above the curves represent the situation where p exceeds break-even and the decision maker should build big.

The region below the curves represents combinations of p and q for which the decision maker should build small. This region can itself be split into two parts according to whether q is greater or less than 0.78. If greater, then the decision maker should expand if the initial demand is high. Otherwise he should not expand.

Notice that the point q = 6/7, p = 7/10, at which the probabilities take their base values, is well within the "Build Big" region and so small changes in these probabilities would not alter the decision to build big.