

[SQUEAKING]

[RUSTLING]

[CLICKING]

RICHARD DE NEUFVILLE: So the reality of design analysis is that the inputs and the context-- that is, demand for whatever the system is producing or delivering-- are uncertain. They have distributions. We don't really know, we can't predict exactly how the system will perform, because we don't exactly know what the costs are.

We don't know what the demands are. We don't know other things about it. So we won't be able to predict it exactly.

So the performance of a system, when we calculate it, has to be some kind of distribution, some kind of range of possibility from the lowest to the best, with different possibilities of the different points along the distribution. And it's important to realize that the outcomes of a system are not the same as the inputs to the system. That is, if you put in a normal distribution, you don't necessarily get a normal distribution out. You're almost certain not to get a normal distribution out.

Why is this? It's because we have complex systems. They perform nonlinearly. There'll be capacity constraints so that you can get the downside, but not the upside.

There's all kinds of reasons why what goes in gets mixed up in terms of the system. And you get a different kind of distribution outside afterwards. So keep that in mind.

Now, let me give you a generic view of the kind of distribution that can occur. And let's see what we're going to be doing about it. So in general, we have a distribution for the design outcomes, represented here in terms of a probability distribution function, not a portable document.

So we have the possibility that things don't go so well, that there are losses, there's poor performance some way. That's on the left-hand side of the graph. And we have some medium, perhaps. It doesn't have to be unimodal, but there's various high points, and more likely situations.

And then, we have the possible upsides that may occur if everything goes really well and if people really like the product that you might be managing, and so forth. So there's some kind of distribution. So what do we do about it?

Well, what we want to do is to have a proactive design. That is, the objective is that we don't have to accept our distribution that may come from the base case design. We can say, right, if this is what's happening, what can we do to improve the design?

So what can we do proactively is first of all, to reduce the downside risks. How could you make it so that those downside risks that you anticipated in your first design don't occur? Well, one way to do it-- not the only way-- is to say, we're going to build the system smaller. So if it's built smaller and it's not succeeding, then you can lose less.

So that's one way to do it. But there are other ways. But in any case, we can think about ways to reduce the downside.

Complementarily, we can take advantage of opportunities. So for example, if the product does really well and if there's a high demand for it, we may be able to expand our facilities, our system so that we can rapidly take advantage of it. We've developed a new gizmo or toy in advance of the Christmas season. It's very popular. How do we go through and ramp up the production and sales to take advantage of it, for example?

And at the same time, what we'd like to do is certainly reduce the overall cost-- at least, the beginning costs-- so that we have the win of avoiding the downside, the win of taking advantage of the upside at a lesser cost. So what do we then do? So here's what I've been suggesting-- the proactive design shifts the distribution right.

So if you had your base case design, which was like this, that somehow you have gone through things, which shift the distribution towards the positive side, both by cutting the downside tail and by extending the upside tail. And you do that by exploring the various ranges of possible things that you could do, the flexibilities that might help you do that. So we want to end up, then, with the win-win solution at lower cost, less loss, more gain, less investment compared the alternatives, at least at the start.

In a graph form, what we have here is the curve that we have. We try to chop the downside, do things that reduce the possibilities that we have really big losses, while at the same time being able to extend the upside if things go really well. So that is the essence of what flexibility design attempts to do. And it very often succeeds, thank goodness.