

The following content is provided under a Creative Commons license. Your support will help MIT OpenCourseWare continue to offer high quality educational resources for free. To make a donation or to view additional materials from hundreds of MIT courses, visit MIT OpenCourseWare at ocw.mit.edu.

PROFESSOR: OK, so let's situate ourselves where we are. So we're starting to really make a lot of progress here, moving in the V model. Today's topic is session five: tradespace exploration, concept selection, and PDR-- preliminary design review.

So there's a lot to talk about, and I'm going to go relatively quickly through this. So, first I want to talk about decision analysis. Fundamentally, when you come up for PDR, you have to make a big decision which is, what concept are you going for? What is your system architecture? That's a big decision. You don't have all the details yet. The design is not fully done, but you've chosen your key architecture. That's a big decision, and there's a whole branch of research and science called decision analysis. So I want to tell you about that.

Talk about some of the issues in concept selection, and then give you some tools for doing that in a kind of organized way. There's a couple of-- what I would say, simple methods Pugh matrix and multi-attribute utility-- that are relatively, I think, straightforward. And then there's a bit more advanced concept called non-dominance-- Pareto frontiers, getting into multi-objective optimization.

And so, at a minimum, when you choose your concept, you want to choose a non-dominated concept. And I'll explain what that means. And then we'll close with discussing what is a PDR? What do you do with the PDR? What's the purpose of it?

So, here is another way to explain essentially the flow of things as we come up to PDR. So the way you can think of it as, you start with a need, right? And your requirements. That's-- we've talked about that before. And then you have this creative phase, which we talked about last time, right? You come up with a lot of ideas, a lot of concepts, and the funnel that you see here is going up. So each of these shapes here-- the circle, the triangle-- they're meant to represent fundamentally different concepts, different architecture, different way of satisfying the requirements.

And so there's a lot of different things you can do, and the funnel opens up. This is sort of the size of options or alternatives you're considering. And then you have to filter this down. In the end, you can only build one system, right? You have to choose one architecture. And this filtering down happens typically in two steps. One is what we call concept screening, which is

more qualitative. And then the actual concept selection is more quantitative, based on models, based on data, based on metrics. And then you arrive at PDR.

So in this sort of simplified view, we've chosen the square as our concept. And we've made that decision by looking at these different metrics. So you can see that on this graphic here, we have metric one, metric two-- you can see the triangle is dominated, right? The triangle doesn't do as well if we want to maximize metric one and two, so it wouldn't make sense to choose the triangle. But the circle and the square-- circle does better on metric two, but less well on metric one. The square does better on metric one, not quite as well-- and they crossover at some point.

So, whatever-- how you weigh this in decision making, the decision is made. We're going to go for the square. That's our concept. And then you move to the actual system design, post-PDR, which is all the details. So we're expanding the box now, the square, and designing all the details within-- the structure, the avionics, the software, the controls. And that's design modeling and optimization.

And so here I'm showing you metric three, which is maybe another metric, and then design variable x_1 , which is some pretty detailed decision that you make within that concept, and you're looking for the optimal, in this case, minimizing metric three. It's this point here. And then that gives you your-- and I put this deliberately in quotes-- optimal design. That's what you show at the CDR, and the critical design review, and that's what you go and build.

So that's roughly the flow of decision making. Any question about that? That general flow? So what we'll talk about today is essentially this portion here, concept screening and concept selection. So, since you make a decision, what is decision analysis?

So decision analysis, in general, is methods and tools for ranking and choosing among competing alternatives or courses of action. And I just want to make-- the word alternative. A lot of people use the word option and alternative sort of as synonyms. They're not. OK? They're not. Options and alternatives are not the same thing. An alternative means if I choose A, I can't choose B and C. But the fact that you've chosen A, or the square, or whatever your concept is, means that you cannot choose the other ones.

An option means, well given that I've chosen A, I could also add something on top of A, right? It's kind of like if you were buying a new car, are you going to get the winter package or not? That's an option, but you've chosen a particular vehicle. So, that's important. So alternatives

are mutually exclusive choices or courses of action. So you need the alternatives, those are the components-- four components of a decision. You need criteria, which means how are you going to evaluate and compare the alternatives?

Then you need value judgments. You now apply the criteria to the alternatives to compare them, and then finally you need a decision maker, which is either an individual or a group with preferences to then make the choice. Those are the four ingredients of decision analysis. And if we show this somewhat graphically-- so here's our alternatives. We can do A, B, or C. The criteria in this case are cost, performance, and risk. Then we basically evaluate, in this case, option or alternative A, I should say, for cost, for performance, for risk. And this gives us some score or number, and then we need to figure out a way to combine this to rank the alternatives.

And when it comes to ranking, there's a distinction between what's known as an ordinal scale and a cardinal scale. So in an ordinal ranking all you really care about and do is say, what's the best choice, the second best choice, and the third choice. You ordered the alternative. But you don't know in an ordinal ranking or ordinal scale, whether one and two are really close, and then there's a big gap between two and three. You don't know that, and you don't really care. All you want is the rank order.

When you put things on a cardinal scale-- so this is a continuous line-- you actually get, in a sense, both the order but you also see how close the alternatives are. So that's a cardinal scale as opposed to an ordinal scale. OK? So-- yes? Please.

AUDIENCE: How do you get a discrete value for something like risk or something like that? And especially in an ordinal scale, if you have two things that are really close to each other, but the risk is kind of a blurry region, how do you--

PROFESSOR: So then they would be equivalent in terms of that criterion. And you can actually have ties, right? You can have alternatives that are tied. And I'll talk about-- this is one of the issues in concept selection is, how do you do tie breaking? Good question. Any questions at EPFL? I know this is very abstract. But fundamentally, that's what we do. Any questions over there? Is it clear?

AUDIENCE: Yes.

PROFESSOR: OK. Good. So let's keep moving. So what are the issues? I think that's-- so multiple criteria,

how do you deal with them? We usually have multiple criteria and how do we-- and eventually, when you make a decision, when you pick a concept, you're somehow combining these criteria into a single measure, whether you're doing that explicitly or not. So how do you deal with these multiple criteria? We'll talk about that.

Here's the point that you brought up. What if there's ties? How do you break them? Group decision making versus individual decision making-- who gets to really make the decision? Right? And this relates to the stakeholder analysis. Who has the power to make the decision in the end?

And then uncertainty-- did we choose the right criteria? Did we evaluate each option properly? And then the big one is are the the best alternatives represented? So you can choose among A, B, and C, but maybe there's this much, much better architecture concept D out there, but it's not in the decision set. Because it didn't-- nobody thought of it, or some other thing happened, and it didn't come up in concept generation. Right? So there's a lot of uncertainties when you're doing this. So you have to be aware of it.

So let me talk about two-- again, quote, end quote-- "simple methods" for doing concept selection. The first one is called Pugh matrix, and then I'll talk about utility analysis. So the Pugh matrix, essentially, is a very discrete method. It uses essentially plus, zero, and minus to score alternatives relative to a datum. A datum is a reference decision, a reference architecture. And this is named after Stuart Pugh, who is a British professor of design, essentially. And it's a very simple method. And it's used extensively. It does have some pitfalls, as well, but it's very well known. So Pugh matrix

And then the second one is called utility analysis. So it essentially maps all the criteria to a dimensionless utility, this is between zero and one. So zero is a completely useless system. It has zero utility. And one is a system that has perfect utility. It can't really be improved. You could try to improve it but your satisfaction would not increase, OK? And this has very deep mathematical theory, particularly the Von Neumann and Morgenstern utility that was basically developed in the mid-twentieth century.

Anybody seen the movie *A Beautiful Mind*? Yeah? What is *A Beautiful Mind* about?

AUDIENCE:

It's about John-- mathematician John Nash who's a mathematician who won-- I forget exactly what prize he had won.

PROFESSOR: He won the Nobel Prize.

AUDIENCE: Was it the Nobel Prize? OK.

AUDIENCE: Founded game theory.

PROFESSOR: For what? For game theory, right. And so one of Nash's contributions was the idea of a Nash equilibrium. So we have multiple decision makers in gaming, and the question is-- you have your current strategy, but is there a better move you could make? And if you could make a better move that improves your satisfaction, or your utility, but by doing this, you decrease somebody else's utility, you may not go there. So if you reach a Nash equilibrium, that means that everybody believes that this is the best they can achieve.

And how do you measure that? Well Nash actually developed his Nash equilibrium and game theory building upon theories of Von Neumann and Morgenstern. So there's a very deep theory behind it. I'm not going to go very deep into it, but so you know that there's a lot behind utility analysis.

OK, so let me just talk you through the basic steps for doing a Pugh matrix analysis. And I'll show you a very simple example. First step, you choose or develop the criteria for comparison. So, what's important? How are we going to decide this?

Selection. And obviously this should be based on a set of system requirements and goals. Now, there are two flavors of requirements. Remember, what-- let's ask here at EPFL. You remember the two flavors of requirements? What were the two flavors of requirements?

AUDIENCE: [INAUDIBLE]

PROFESSOR: I mean, there's actually six types, right? There's functional requirements. There is constraints. There's interface requirements, and so forth. What I'm trying to get at is there are requirements that use shall statements and then there's requirements that use should statements, right? So the shall statements are what? What do the shall statements imply?

AUDIENCE: So basically, for the shall, it's compulsory, and for should, it's a goal of the system.

PROFESSOR: That's right. Exactly. So, in some sense, the shall state-- the criteria associated with shall statements don't necessarily make the best criteria for the Pugh matrix, because by definition, the concepts that you're evaluating should all satisfy the shall statements. Because if they

don't, they're essentially infeasible, right? And you shouldn't select among them in the first place.

So architectures, or concepts, or alternatives that violate shall requirements do not satisfy the must have requirements. And you shouldn't-- so you shouldn't use criteria that are associated with hard constraints. That's what I'm trying to say.

Now the requirements that are associated with should statements, those are good criteria for this Pugh matrix, because the degree to which you satisfy this, the goals are somewhat variable. Therefore, you can then compare amongst alternatives. OK?

So that's number one. Number two-- thank you, that was good-- select the alternatives to be compared. So that those are coming out of concept generation and one little bit tricky aspect here is when you're putting together the alternatives, the concepts to be selected from, they should be represented at a similar level of detail of abstraction. It is not good to have one concept that's really detailed, and you know a lot about it, and often it's kind of an existing system or something you've done before, and then there's this other concept that is very fuzzy.

Why shouldn't you compare something very detailed with something very, very poorly defined? Why do you think that's not a good practice? What do you guys think here? Yes.

AUDIENCE: Is it like-- it's not fair, because--

PROFESSOR: That's right. It's a fairness issue, and because the concepts that are poorly defined often, you're too optimistic about them. That's sort of the typical-- this looks really good, very promising, but it looks very promising because it's kind of ill-defined. So that's important. Try to represent all the concepts at about the same level of detail when you make the comparison. That's number two.

Number three-- you actually go through and generate the scores. And the key thing here-- this is maybe the most important thing about the Pugh matrix is you always use a datum. The datum is a reference-- one of the alternatives that you're going to compare is your reference. So what you do is you always compare every other alternative in the set against the datum, and you say for each criterion is it better? Is it about the same, or is it worse than the reference, the datum?

So you don't compare all the alternatives pairwise against each other. You always compare

against the datum. Then you compute the total scores and I'll show you this in the example, and then there's variations like-- rather than better, equal, or worse, you could go to a five-point. Much better, a little better, or about the same. So there are some variations on the scoring, but the classic method just uses better, the same, or worse.

So here's some--

AUDIENCE: Excuse me, just a question.

PROFESSOR: Yeah, go ahead.

AUDIENCE: Just about the number three, about generic scores and user datum-- by doing so, don't you exclude a solution that could have, under the generic picture, be better than others?

PROFESSOR: What do you mean by general picture?

AUDIENCE: You should take all the solution that you have, and at the end, one performs slightly better than the other, but you couldn't spot it by using the datum?

PROFESSOR: Yeah. So it's a very good question. And so this is a question of resolution right? How good is your resolution? The main purpose of the Pugh matrix is not for making the final choice right away. The main purpose of the Pugh matrix is actually to eliminate concepts that are-- you can tell already that they're not competitive. And you don't typically do Pugh matrix one time and then say that's the best one. You use it, essentially, to eliminate purely inferior solutions. Does that make sense?

AUDIENCE: I'm really sorry but the video cut just during your explanation.

PROFESSOR: OK, so let me repeat this. So the purpose of the Pugh matrix is not to select the final concept as a result of doing the Pugh matrix, but to identify inferior concepts to eliminate them from the set. OK?

AUDIENCE: OK.

PROFESSOR: All right. So here's a kind of generic example to explain how this works. So here's our Pugh matrix, our evaluation matrix. On the rows, we have our criteria. It could be cost, performance, resilience, effectiveness, and so those are called A, B, C, D, E, and F. So we have six criteria.

And then the column-- each column represents a concept, a design, an architecture. And so

you know we're designing-- I guess we're designing beams here, or levers, or something like that. So this is a solid cylinder. This is a cylinder that's hollow. This is a square, it's cross-sectioned. This is a triangular cross-section. You get the picture, right? There they're qualitatively different.

And then you note which one is our is our reference here? Which concept? Number seven, right? If you look at seven this sort of shallow triangle. Seven, it says datum. You see that? Which means that this is our baseline, our reference, our datum.

And then what you do is, as you fill this in-- and you can do it two ways. You can do-- pick, say, concept one, the solid cylinder, and then compare it against seven for all criteria. So A, so concept one is better than concept seven in criterion A. So it gets a plus. It's also better in B. It gets a plus.

It's worse in C and D, so this is a minus compared to seven, and so forth. So the two ways of filling it in is you do concept by concept, column-wise, or you can do it row-wise.

And there's some arguments in favor of one or the other. But essentially, you fill in this matrix by always comparing against number seven. Alexis, what's your question?

AUDIENCE: Yeah, just a question. Shouldn't we weight the criteria? I mean, for example, if the criteria A is more important, shall we put more weight on it?

PROFESSOR: No. No. At this point, in the Pugh matrix, the criteria are unweighted. This is very important. You don't say B is twice as important as C. They're unweighted when you first do this. Eventually, when you do the final selection, you will probably weigh them. But when you first do this, there's no weighting. They're equal. OK? Yes? Go ahead.

AUDIENCE: Are there any rules in picking your datum or do you just pick arbitrarily?

PROFESSOR: So I'll get into that. It should be a design that you think is competitive, that [INAUDIBLE] maybe you know it, maybe it's an existing sort of system, there's data on it, a known quantity. But if you don't have if you don't have any information, if they're all equal, then it's a random choice. OK?

So, then once you've filled in this matrix-- so S here is for same, right? About the same. Then you essentially sum all the pluses, you sum all the minuses, and then you have the net result, unweighted. Right? So, here you can essentially compare that, you can see that concept one

is better in three and worse in three than the datum. You can see that concept two is inferior in three criteria, better in two, and the same in one. But we're not actually summing. We're not actually summing at this point, because we're not weighting these, OK?

But it gives you a basis of understanding. OK so let's do a quick partner exercise. What do you see as the main advantages and potential disadvantages or pitfalls of this method? So turn to your partner, discuss this for a few minutes, and then we'll see what you came up with.

[MICROPHONE NOISE]

AUDIENCE: [INAUDIBLE]

PROFESSOR: All right. Good. So let's see advantages. Let's get a couple here, and then we'll sort of go back and forth between here and EPFL. So what are-- what's good about this method? Yes, go ahead.

AUDIENCE: So, Nate and I were talking that it is a relatively simple decision making process. So instead of just going with some sort of gut feeling, or not looking at everything, you can get stuff down on paper and actually see these criteria and how they relate to each other.

PROFESSOR: Yeah. So simplicity. What else? Go ahead, Sam.

AUDIENCE: It allows you to look at a large array of concepts very quickly.

PROFESSOR: OK large samples. Sorry. Quick, right? You said quick. . OK. What advantages, EPFL, what did you guys come up with? So simple, it allows us to look at a lot of concepts. It's quick. What else?

AUDIENCE: [INAUDIBLE] we got [INAUDIBLE] the same things. This method is simple, and it's pretty fast and pretty straightforward, too. So this is a good way to get rid of some of the concepts pretty fast.

PROFESSOR: OK. Anything else that we didn't [INAUDIBLE]? There's one more that I would like to put on this list. Veronica, go ahead.

AUDIENCE: [INAUDIBLE]

PROFESSOR: Qualitative.

AUDIENCE: [INAUDIBLE] I kind of like [INAUDIBLE] improvements, rather than this is better because the

requirement is--

AUDIENCE: Please put on the microphone.

PROFESSOR: I think-- can you just quickly repeat that?

AUDIENCE: Just that it's qualitative.

PROFESSOR: And that you see that as an advantage, actually.

AUDIENCE: I think you can see it both ways.

PROFESSOR: Yeah. So I think what you implied with is it stimulates the discussion, right? Stimulates debate.

AUDIENCE: There may be another criterion on the positive side.

PROFESSOR: Go ahead.

AUDIENCE: And it would be that it's extremely easy to explain to [INAUDIBLE], to [INAUDIBLE], to whoever your customer is, because-- and they can relate to it. If the criteria are mostly physical or scientific, I mean, larger, stronger, lighter, cheaper, it's probably very intuitive, yes, as I say.

PROFESSOR: OK, great. Now, downsides. What are the downsides of the method? Go ahead. Make sure you--

AUDIENCE: I think it definitely depends on which you pick as a datum. Like, if you pick the clearly best one as the datum, then it's not going to give you much information. Or if you pick the worst one as the datum, then all of them will be better. And you kind of don't really gain much information from that. So you have to like iterate with different datums.

PROFESSOR: Yes. There's actually research on this. There's some research-- this matrix, this Pugh matrix method, by the way, has been studied scientifically quite a bit. There's quite a literature on it. And the point that you brought up was studied. They would basically give the same criteria and the same alternatives to different groups of people, but then give them a different datum. And it's a little tricky-- is the result different because it's a different group dynamic? Or is it-- but this has been done statistically, and the choice of datum has an impact on the outcome. Very good.

That's actually a pretty subtle point. That was excellent. What are some other-- yes, please go ahead.

AUDIENCE: While it's very easy to implement and anyone can do it, it is very subjective. So if you have somebody that doesn't have a whole lot of experience in a certain area, they may get a completely different answer than somebody who is an expert looking at the same type of criteria.

PROFESSOR: Subjectivity-- which basically, repeatability-- may be low, right? You want to have robustness of the methods such that if you repeated this multiple times or you gave it to different groups of people, you want to have some confidence that the results will come up similarly. Very good. Let's see at EPFL, any downsides? So datum dependent, subjective-- what else?

AUDIENCE: So what we are thinking is that the criteria depend, from group to group. It doesn't seem to be clearly defined what you need. Although, there are some that are straightforward and our other point--

AUDIENCE: Datum dependence? That's what we--

AUDIENCE: Yeah, datum dependence. If it was a very easy solution, or if a very bad solution as a reference, then all the others look good. And finally, if you have a-- if you judge on the overall concept, and maybe you discard one or two solutions, but inside those concepts you have ideas that you could keep or use, then they might be lost.

PROFESSOR: OK, so loss of sub-concepts. OK. Good. Anything else here? Mike.

AUDIENCE: We were talking about how it's very easy to tune your metrics to be-- like to be biased.

PROFESSOR: I see. Such gaming, right? Gaming-- gaming the-- basically producing a matrix to give you the answer you want, that you already had preordained, sort of. Right? That essentially what you're saying. . OK. Very good. So these are all-- go ahead.

AUDIENCE: Just sort of one last disadvantage of the method that I think is a bit interesting, because I think the way of grading everything with plus or minus can be [INAUDIBLE], because you can imagine a system, for instance, that is very, very good in one criteria, and then pretty bad in all the others, and it could be a good concept that is disregarded nonetheless.

PROFESSOR: What's your name?

AUDIENCE: Bastian.

PROFESSOR: Pasqual?

AUDIENCE: Sorry?

PROFESSOR: Pasqual. Is your name Pasqual?

AUDIENCE: Bastian.

PROFESSOR: [INAUDIBLE]?

AUDIENCE: Bastian.

PROFESSOR: OK. Keep that point until-- we're going to talk about exactly that point when it comes to non-dominance, OK? So please reserve-- this is great. Just, we'll come back to that in a few minutes, OK? All right, so this is great. I think you really got it. This is exact-- this is a very good list.

So, you may say well, OK this is good for beams and you know very simple things. So, I want to show you an example of an application of the Pugh matrix to what I would argue is a fairly complex architectural decision. And what I'll talk to you about is a thesis that was done here at MIT in 2003, so a while ago, 12 years ago, by Brian Smith. He was an STM student, and was one of the easiest students ever to advise. Very knowledgeable in nuclear propulsion and nuclear power.

He's now the branch chief for nuclear power propulsion at NASA Glenn research center in Ohio. And so, at the time, NASA had a mission that was kind of high priority called JIMO-- Jupiter icy moons mission. And when you go out to Jupiter, and you want to not just do a fly by but you actually want to go to different orbits, you want to do high power imaging, you need a lot of power.

You're far from the sun. RTG's only give you 100 watts or so, per RTG. So nuclear power is pretty much the way to go. But there's a lot of different choices of different nuclear reactor architectures that would potentially-- and then couple that with the propulsion system.

So that was the challenge was sifting through the large space of possibility for nuclear electric power, nuclear exploration class electric power and propulsion systems. And you know, here's the way the study was set up is the expansion phase, the filtering, and then the screening phase. So here's the architectural space the way it was defined.

So nuclear electric power and propulsion is what we're after. And there's different ways to generate these alternatives. So the most important is the design vector, the type of reactor, the operating temperature, the power conversion scheme, the heat exchange, and then the fuel type. And then there were requirements. So this is the power range, the delivery timeline, the fact that you should be able to do it in a single launch, and then the operational lifetime of the system. So those are the shall requirements. Every architecture had to meet these requirements otherwise it would not be considered.

And then there were some things that were assumed as constants. They're shown here. There was a policy vector in terms of funding profiles, international partnerships, at what altitude you could insert the system into orbit, and then influence on future missions. And the objective vector is the criteria for the decision.

Technology-- TRL is Technology Readiness Level. How much infrastructure do you need? Complexity. Strategic value to the nation. This is a little fuzzy, but basically what this means is, could you use it for other applications than exploring Jupiter? The schedule. The launch packaging. Power. Specific mass, which is watts per kilogram. The lifetime.

So it's interesting you have lifetime here and here. So there's a minimum lifetime, right, and if you satisfy this, the extra life you get can actually be used as a decision criteria. How does it interact with the payload? Do you need a lot of shielding? And then adaptability? Can you adapt it for different missions? Yeah?

AUDIENCE: In the previous slide, it listed both possible-- on the previous slide, we had both possible and feasible on there as different kind of criteria. What do you have for a difference for those?

PROFESSOR: So essentially, possible is, this is sort of the combinatorial space. And then, applying the hard constraints based on the requirements vector gets you from possible to feasible. They satisfy your must have requirements, your shall requirements, and then you use, in screening, you use the objective vector-- these criteria-- you use to go from here to here.

So let me show you what this looks like, and I'll post this thesis, if you're interested. It's really-- the whole thesis is about doing this process. So here we have-- each column represents a different reactor architecture. So you can see in the legend what these things mean.

So for the reactor type, we have a liquid metal, gas cooled, or heat pipe type reactor. These are three types of nuclear fission reactors. For fuel, we have two types of fuel, if I remember

correctly. We have UO₂ or UN. For the temperature at which the reactor would be run, we have medium or high. For the conversion-- this is the thermodynamics. How do you get the heat out of the reactor? We have a Brayton or a Rankine cycle, or a thermoelectric reactor, which is basically directly converting the heat to electricity without a working fluid. And then, D is the heat exchange architecture, which is either direct or indirect.

And so you can see that each of these columns represents a different kind of architecture, and they're colored here by-- the dark is a direct heat pipe with liquid metal reactor, and so forth. So that's essentially-- what you're seeing here is the filtered set that is deemed to be feasible. Potentially feasible.

And then the actual screening is done based on this. So this is the actual Pugh matrix here. So we have the concept combinations, which are the columns. The datum here is the one that's shaded in gray. So it's a liquid metal reactor with a thermal-- it's a liquid metal thermoelectric reactor with indirect heat exchange, UN fuel type and a high operating temperature. And why was this chosen?

Well if you look here, it says SP100 reference. So the US actually has launched-- at least, officially, as far as we know-- one nuclear reactor into space, and that's the SP100. You can look it up. There's a history behind it. And that's the architecture of the SP100. So it was launched. We had data on it. It's a known quantity.

The Russians, it turns out-- Russia has launched a lot of nuclear reactors over the years. But this is sort of the US baseline, if you want to call it that. And then all of these other combinations are compared against the SP100, the reference reactor.

And so, you can see here that the comparison is drawn again. So zero means it's about-- so for TRL, it means it's about equally mature. A plus means it's actually more mature. A negative means it's less mature than the SP100. And at the end of this, you can actually look at the total pluses and zeros and minuses, and then in this case, what Brian did, he did calculate a net score. That's an unweighted score.

Basically, if you give one point-- so let's look at the first one. So it has four pluses, right? It has four pluses, and it has five equal to the SP100, to the datum, and has two minus. So it's worse than two, better in four criteria, and the same in five. So the net score would be plus two, because it's better in two criteria, net. But it's unweighted.

And so in order to figure out what-- based on this net score, so a zero net score would say it's about equal to the datum, but it maybe in other, different criteria. But essentially it's about the same. If it's a plus two, then it is potentially better than the SP100. But you don't know that 100%, because you haven't weighted the criteria. It's really used for screening.

So the ones that are circled are better, essentially, than the datum. And the ones that are-- this is-- I'm sorry, the rank. This is the rank among the set. So if you have a one here, you're in the rank one. And there's two architectures, the first two are rank one.

And then we have this architecture, which is ranked two. The datum itself, plus two other architectures, are rank three. And then you see that there's some-- like, for example, this architecture here is rank eight, right? So it-- why is that? Because it's not better in any criteria than the datum. It's the same in six. And it's worse in five, right? So that gives you an indication that this particular concept probably is not very attractive, and you would eliminate that. OK?

So that's essentially an application of the Pugh matrix to a pretty real-world, complex problem. So what the Pugh matrix is for is essentially structured-- structuring and representing the evaluation procedure. It serves as a common visual. It provides some amount of discipline, and helps break down self-sealing behavior, meaning-- what that means is people defending their concepts without sort of seeing the bigger picture. Encourages teamwork. It helps you eliminate weaker ideas, retain a set of stronger concepts. And then this is interesting-- divergence. It helps to identify opportunities for combination.

There's also a way to use this Pugh matrix methods multiple times. And the way you would do this is you would say, OK so these are the stronger concepts here, right? The ones that do equally well or maybe even a little better than the datum. What are the strengths and weaknesses of each concept and can we hybridize them? Can we create some hybrid concepts?

And so you eliminate the weaker ones. The surviving concepts are kept, and then maybe hybridized with each other. So you expand and create more concepts again. Maybe a little bit more detailed. And then you apply the Pugh matrix again. So it's iterative application of the Pugh matrix, where you eliminate, you keep strong concepts, you hybridize them, and then you repeat it two or three times.

And there's actually research on this, as well. Professor Dan Frye, who is here at MIT in

mechanical engineering, has done research on that and shown that that can actually lead to very good outcomes. So this sort of gradual-- so then, you don't have just one expansion and contraction of your concept space, but it sort of goes multiple times. But eventually it does converge. So this is a more refined application of the method.

What Pugh matrix is not good for is automatic decision-making, right? Completely controlling the process. So, you know, the idea of automating the Pugh matrix-- it should not be automatic. And it's not really done-- it's not really good for trade studies.

Challenges-- [LAUGHS] so here's some quotes. "People who have a lot of experience exhibit an impatience. Get on with it. This procedure holds us back." Right? "Strong-willed individuals who have a lot of experience and whose initial concepts have not emerged in the final selection commence a defense based on emotion, experience, and bluster." Right?

So there's this social dynamics that gets unfolded here. So, therefore, it is recommended to do Pugh matrix with a facilitator. Somebody who-- similar to brainstorming, right, where you have a facilitator. So you have a facilitator for Pugh matrix. So somebody who controls the flow and pace of the session, records the results, tries to maintain some discipline. Always compare against the datum.

I've seen Pugh matrix sessions where you start with the datum, and then by the third or fourth concept, people are starting to do pairwise comparisons, and the datum is sort of lost. That should not happen, right? Preventing tangents, but encourage clarification. What do we really mean by criteria? Clarification of the concepts, and then opportunities for divergence in hybrids.

So the critique of the method is-- I think a lot of these things you've already mentioned. The ranking depends on the choice of datum. The weighting. You know, there's no real weighting in the classic method. However, you can implement a weighted version of the Pugh method, but then you have to have the whole discussion about the weightings, *a priori*.

Next, I'm going to talk about multi-attribute utility. And so Pugh matrix and multi-attribute utility could give you a different rank order of alternatives. So how do you reconcile that? The most important criteria may be intangible and missing from the list.

And my personal opinion on this is that the Pugh matrix is useful and simple to use. It stimulates discussion about the criteria, the alternatives, but it shouldn't be your only means of

concept selection. But it's usually one of the first things you should do because you learn a lot from doing it. OK?

Any questions about matrix before I move on? Yes, Sam?

AUDIENCE: Should it always be done with a group that's working together and discussing the concepts? Or could you also have people individually filling them out and then combining their results?

PROFESSOR: So yeah, usually you do it as a group because in the discussion with the group, there's a lot of clarification that happens. Now if you do it separately, and you firewall people from each other and have people do this individually and then try to combine the results later, that's actually called-- that's closer to a method called the Delphi method. Which I'm not going to talk about here.

But the Delphi method means you're asking these kinds of questions to experts. They give you the answers, and then you anonymize the answers and reflect the combined answers back to the group. And then you eliminate the social dynamics, and you can see whether was my assessment an outlier, am I close to the center of the group? So the Delphi method has a lot of benefits, but it's not usually-- you don't have the social dynamic in the same way, which is sort of essential to make this work. Good question though. Any questions at EPFL about Pugh matrix before I move on?

AUDIENCE: No, it's clear.

PROFESSOR: OK. Go ahead.

AUDIENCE: It is useful to almost have an independent group do this, as opposed to everyone generates their ideas and then everyone's trying to fill out the matrix so that their idea wins.

PROFESSOR: Usually not, because if you give it to an independent group, they'd know much less about the problem. They probably weren't involved in generating the alternatives, they weren't involved in selecting the criteria. They're probably not as knowledgeable. And so there is an advantage of doing it within the team. But then, you know, before you actually go and pull the trigger on the final-- that's the PDR. You get vetted by independent people. That's what happens at the PDR. OK.

AUDIENCE: One point here is that you can invite external people to join your discussion. So you have your design team. And it's very good to have the experts to come, sit in the meeting, and then give

their opinion, and with the Pugh matrix, as it's pretty straightforward and simple, you probably shouldn't then bias the result, but you could probably enrich it.

PROFESSOR: That's a great point. I agree with that, [INAUDIBLE]. Absolutely. Yeah.

OK. So let's switch gears to utility theory or multi-attribute utility theory. So what's that all about? So like I said, utility is a very deeply rooted concept in economics.

And so it's defined as, "Utility is a measure of relative happiness, or satisfaction, or gratification gained by consuming different bundles of goods and services." This is sort of economic-- very-- this is how economists talk. So it's-- the idea is you want to choose a concept or alternative that will maximize your happiness, your satisfaction, your gratification.

And whenever you buy something, when you go to Anna's Taqueria, or wherever you guys have lunch, or at EPFL you go to the Rowlock center, and you pick from the menu, you're actually doing just that, right?

You have-- how much money do I have in my pocket? How hungry am I? Am I a vegetarian or not? That would filter out some menu items right away, right? And you pick, every day. You pick your lunch. Well, you're doing exactly this. You don't really know that you're doing it, but in your mind, as you're picking from the menu, and then making your choice, you're maximizing your utility at that moment. So we do this every day.

Here, in this class, we talk about designing complex systems. So we need to do it a little bit more deliberately, but this is not just some abstract thing. We do this every day.

And so the idea is, essentially, we have this Consumption Set X , which are our alternatives-- mutually exclusive alternatives-- and then we map that to the real scale. We rank each member of the Consumption Set. So we map the alternatives on this utility scale. In this case, the cardinal scale between zero and one.

So, the way this is done, in order to do the mapping from the criteria to utility, you need these mapping functions. And we call these utility functions. They're called utility functions. So basically we have-- I'm going to use J here for-- J means, essentially, your attribute, your objective. And then U is your utility.

And so there are different shapes of utility functions. And different scholars call them differently. So this is Cook. He says, this is smaller is better, or larger is better. And you see

they're essentially monotonically increasing or decreasing curves. Messac, another scholar in multi objective design, calls this class 1S, class 2S. So what would be an example of a smaller is better utility function? What would be an example of that? Go ahead.

AUDIENCE: An example would be the weight of a launch vehicle. It's decreasing.

PROFESSOR: The launch mass?

AUDIENCE: Yes.

PROFESSOR: OK. Launch mass. OK. What about EPFL? Smaller is better.

AUDIENCE: Cost.

PROFESSOR: Cost, right. So, larger is better?

AUDIENCE: Revenue.

PROFESSOR: Revenue. Maybe range, endurance, reliability, right? Then there's-- the next one is strictly concave or strictly convex, which is nominal is better. So the idea there is there's a sweet spot, right? You want the performance or the attribute to be a pretty specific value, and if you deviate from that on the up or down side then the utility decreases. So what would be an example of that?

AUDIENCE: Size of a meal.

PROFESSOR: Size of a meal, OK? Well, some restaurants may beg to differ, right? All you can eat.
[LAUGHS]

AUDIENCE: [INAUDIBLE] what their objective is.

PROFESSOR: Yes. All right. Very good. Yes, Sam?

AUDIENCE: [INAUDIBLE] has to interface with something else [INAUDIBLE] you're not [INAUDIBLE] it [INAUDIBLE].

PROFESSOR: Yup, so if there is a very specific interface condition. Yeah. OK, go ahead at EPFL. Nominal is better.

AUDIENCE: Like, for example, ambient temperature.

PROFESSOR: Yeah. So temperature-- ambient temperature, right? Humans are pretty-- we have a fairly narrow range where we say, this is good. Now we can put on a sweater or-- but it's a fairly narrow range.

So, great example. Then the next one is range is better, which is concave or convex. So this is the idea that as long as you're within this interval, within the interval itself, you're sort of indifferent. Right? But then when you drop outside the interval, then utility decreases or increases. And then the last one is very exotic-- non-monotonic utility functions that have multiple peaks. They exist in theory, but in practice, you almost never see them. I can't give you a good example of multi-modal one. But in theory they do exist.

AUDIENCE: Well, there is one, and it's clearly the landing sites on a planet.

PROFESSOR: OK.

AUDIENCE: Because you have the one injection trajectory, and then you have the primary sites of the primary ellipse, then the secondary. And it's non-monotonic, mostly [INAUDIBLE]. It doesn't go up, it goes down. So you have ideal. Then you are out of it. Then you would land on rocks. And then you have the next plateau, and maybe a third one, on one pass by.

PROFESSOR: That's an interesting-- that's an interesting comment. I'm trying to figure out what's the attribute-- what's the engineering attribute that goes with that?

AUDIENCE: Well, it's the tolerance on the injection. I know that, at JPL-- by the way, thanks for organizing the [INAUDIBLE]-- they made this presentation, also, of the landing of the rovers. And they have the primary ellipse where they try to land-- well, try-- it depends on the entry. And then they have secondary-- if they miss the entry, then they have very bad option. And then, maybe a little bit later, they maybe can land on the next plateau, or in the next crater, which is useful. So they have this non-monotonic, but degrading, function. The first time is best. The second would still work. But in between, there's a part that doesn't go.

PROFESSOR: OK. Got to think about that one. But that's an interesting-- that's a pretty interesting example. So, the main point here is that in order to calculate utility, you need a translation function between the engineering or financial attributes of each alternative and utility.

If it's not-- if it was directly that, you'd just have a linear, right? It would be a 45 degree line, right? Or negative 45. But typically the mapping from your attribute values to the utility is non-linear. And that's what this says.

So then the challenge is well, how do you get these utility functions? Where do they come from? And the answer is you've got to do interviews, you have to survey people-- the decision-makers, the stakeholders we talked about. They're the ones who have these utility functions in their minds, even if they're not-- So you've got to make those explicit.

So here's an example. You may have three different customers for your system. User attribute, some performance attribute. Customer one is shown here. They need a minimum amount of performance, but once you reach this threshold, after that there's no more utility. But then customer two and customer three are different. They want-- their utility increases gradually, and in fact, customer three doesn't see any much utility until you hit this much higher level.

So one of the challenges, then, in designing these utility functions when you have multiple customers is to solicit this-- and there's interview and survey techniques for doing that-- and then combining those. And then the other challenge in combining them is that, remember at the system level, utility is still between zero and one, right?

So if you had, for example, two attributes, and they both give you perfect utility, and you add them together you'd get a utility of what? You have two attributes. Two. But that can't be because the total system utility can never be better than one. So you have to normalize. As you combine the utilities, you have to normalize them, and actually weight them as well.

So here's the equation when you have two utilities. You basically have the utility of the system just from J_1 -- your first attribute-- times the utility of J_2 . This is the mixed term, right, the combined term. Plus the utility of just J_1 plus the utility of just J_2 . Then you have this K factor here-- this capital K factor-- which re-normalizes everything to zero and one. And if you have more than two, then it becomes a matrix calculation.

So it's a little tricky for how to do this properly, but it's well-known how to do this. So the steps in multi-attribute utility analysis are to identify your objectives and attributes, you develop an interview questionnaire, you administer that questionnaire. You then develop your aggregate utility functions, you determine the utility of your alternatives, and then you analyze the results.

And there is one word of caution I want to give you, which is that utility is essentially the surrogate of value, but value often we express in dollars, in monetary terms. But utility is unitless.

So let me just comment about utility maximization. It's very common and generally well accepted. It's a non-linear combination of your criteria-- your decision criteria. The downside of it is the physical meaning is often lost, because the engineering or customer focused attributes all get combined.

So you say this system has a utility 0.83, and this one has a utility 0.75. So one is better than the other, but immediately you want to say well why is that? What does that mean? So then you have to backtrack and reverse engineer how those utilities were calculated.

You need to obtain a mathematical representation for all utility functions, and-- this is a probably not just US-M centric comment, this is probably everywhere in the world more or less-- but the utility function can vary drastically depending on the decision maker. This is a big issue in government programs.

So you guys, it's uniform day, I guess, today. So how long is a typical tour of duty of a program manager in the Pentagon for a big program? What would you say?

AUDIENCE: It's usually three to four years.

PROFESSOR: So three to four years, like it's written here. And what does that mean in practice for these utility functions and for programs? Sorry I'm putting you on the spot here.

AUDIENCE: Well, you don't usually have a lot of overlap between the PM. So, as one goes out, another one's coming in, so you're losing a lot of experience, gaining somebody who doesn't have much experience. So there's a learning curve here.

PROFESSOR: Right. That's right. What I'm getting at is-- so there's definitely that effect-- but the priorities may be different. You know, whereas the prior program manager really valued a lot performance or quick response, and the next program manager-- because maybe the context has changed-- is very, very cost conscious. Really wants a system that is very, very affordable, and is willing to sacrifice performance for that.

So the way you can think of this is the shape of the utility curves for those decision makers has now shifted, and because of that the architecture or the choice of concept that you would go for may be different.

AUDIENCE: This happens to NASA all the time, right?

PROFESSOR: Right. And so this is not just a DOD issue. And in commercial world. You know, a new CEO comes in, new CTO comes in, and these utility curves actually shift. So you've got to be to be aware of this.

So one of the big topics there is-- particularly now with NASA is-- how do you choose an architecture, a concept, that is robust to changing utilities and decision makers. So maybe you don't go for the super duper best, most exciting, most capable concept, but you go for the one that's least likely to be overturned or disrupted with the next administration.

And this is a discussion-- active discussion-- right now at NASA headquarters. You know, what are the investments you can make that are going to provide utility, even if the utility function of the future decision makers changes? It's a big, big topic.

The other thing is, of course, this requires your formulation of preferences. Those K factors, those weightings, *a priori*-- before you've actually scored the alternatives. All right. So the example I want to talk you through is essentially a space tug. So we're going to do trade space exploration. There's a trade space, a design space, of alternatives-- conceptual alternatives.

And we're going to try to understand that using utility theory. And let me explain to you what we mean by space tug. So a space tug is essentially a satellite that has the ability to change its orbital elements-- there are six orbital elements, semi-major axis, eccentricity, inclination, right ascension of the ascending node, and so forth-- of a target satellite by a predefined amount without degrading its functionality in the process.

So here's a picture of the earth. There's two orbits. The space tug is here in the black orbit, and then our target satellite is in the orange orbit. And so the typical process is the space tug waits in its parking orbit-- the black orbit. It gets tasked. It transfers to the other orbit, searches for the target, identifies it rendezvous and approach, docking and capture, does an orbital transfer and then releases the target satellite at the new orbit, verifies its status and then either goes directly to the next target, or returns to the parking orbit.

And as you can imagine, depending on the plane changes that are required here this can be quite expensive. It's less expensive in geosynchronous orbit, and there are some capabilities that we think heard about in the public that US and maybe other countries have as well to do this.

But there's also applications-- commercial applications for it. So, for example, for space debris

removal and things like that. So that's what we mean by space tug.

So what are the attributes? What are the-- so this is based on a paper from 2003 called Understanding the Orbital Transfer Vehicle Trade Space. And so here's three attributes that are combined into utility: total delta V capability-- delta V is the change in velocity. This essentially tells you where can it go? You calculate that from essentially the rocket equation.

The second one is response time. How fast can it get there after it has been tasked? And there there's a big distinction between-- it's almost binary-- between electric propulsion, which is very efficient but slow, and then chemical propulsion. The third is the mass of the observation or grappling equipment. And so this tells you what it can do when it gets there. So the size of target satellites that can be actually interacted with.

Those are the three attributes that define utility, and we combine those into a utility between zero and one. And then in this case, cost or a surrogate for cost is kept separately. This is the vehicle wet and dry mass. These are the cost drivers, and this is calculated from some simple scaling relationships.

And so what we're interested in is looking at the trade space of utility versus cost of the system. So how is utility defined for a space tug? So we have response time, which is bad in electric systems. Total utility is a weighted sum, and then we estimate the cost from wet and dry mass.

So the delta V utility is shown here. This is in terms of meters per second. So two, four, six, eight, 10 kilometers per second. These are pretty large delta Vs. And what's interesting is this is a larger is better, but it's a step curve. It's not a smooth curve, it's a step curve. Why is that?

Because as you reach certain delta Vs, it enables operating just in GEO or a LEO-GEO transfer, or a LEO-GEO return. Right? So you can use the space tug more than once, the more delta V you have. And it's a step curve because as you hit this value of delta V, you can now operate in a different regime.

The capability. The payload mass utility is essentially discretized between low, medium, high, and extreme. So low is for small satellites. Medium satellites up to one metric ton, up to three metric tons, and then more than five metric tons, which are the big satellites in geosynchronous orbit. And then the weighting factors are we're going to capability 30%, delta V 60%, and the time responsiveness of the system 10%. So time responsiveness is not that

important compared to the other two.

So once we have this-- I'm not going to go through all the details of the calculations. That's in that paper, which I'll post. You get this. So this is a cloud of points. Each point here represents a particular alternative or architecture, and we have the cost of the system in millions of dollars. So this is not cheap. All right? This is a billion dollars, two billion, three billion. And then we have this utility, this dimensionless utility. And, of course, what we're particularly interested in is this lower right corner: high utility lower cost systems.

And what we did here is to identify some particularly interesting architecture, as they're shown here in this space. So let me just explain two of them. One of them, here-- and we gave them names that are recognizable. So this point here, it's below 0.4, so it's a rather lower utility system, but it's also relatively affordable. We call it the bi-prop low-earth orbit tender.

So it's intended for use in low-earth orbit and it has a dry mass of about 680 kilograms, a wet mass of 1,400. So the difference is propellant-- about 800 kilograms of propellant-- and has reasonable size and mass fraction. And it's fairly responsive.

Another alternative is what we call the electric geocruiser. Sounds cool, doesn't it? This is basically a space tug that operates only in geosynchronous orbit. So you launch into geosynchronous orbit, but once it's there it can do a lot. But it's electric, so it's kind of slow. Seven hundred eleven kilograms dry mass, 1,100 kilograms wet mass. And it includes return of the tug to a safe orbit. And this is sort of a versatile space tug in the GEO belt.

So that's the idea. You calculate utility, shown on the x-axis, you have the cost of the system. You have this cloud of points. And you start understanding what are the interesting architectures in that trade space. This is the same trade space. But now what we've done here is we've shown the choice of propulsion system is critical. And what you can see here, it's really almost impossible to get to a utility of one. It's very, very hard.

And the reason for this is the rocket equation. You see how these-- so the blue are the bi-propellant architectures. The purple are the cryogenic, so these are [? lux ?] hydrogen propulsion. The yellow ones are the electric propulsion system, and then we also have nuclear propulsion, here, which clearly is challenging and has policy implications. So you can see the nuclear propulsion gets you close to much higher utility, but it's also much, much more expensive. But in all cases, as you're trying to get more and more utility, at some point you hit the wall because of the rocket equation, which is very non-linear.

You know, more fuel-- in order to carry more fuel, you need more dry mass, more fuel and dry mass requires more fuel to push, and the system blows up on you. And you can see this in these curves. So you get some physical intuition in terms of the shape of the space. Yeah?

AUDIENCE: Is there a reason between this layering of cost?

PROFESSOR: Yeah. They're essentially, if I remember correctly, the capability of the system. Remember there's this small, medium, large, you know the size of the grappling equipment? I think that's the tiering that you see.

AUDIENCE: That doesn't change the utility?

PROFESSOR: Well, it depends on the weighting, right? So as you change the weighting among the attributes, that space can get scrambled. This space is valid for that particular weighting. So any questions about-- this is utility theory, multi-attribute utility theory, applied to trade space exploration. Any questions? Go ahead. EPFL, do you have a question?

OK. So now what you see here is implicitly, is there's nothing here in the lower right corner. There is no system that has perfect utility and is low cost. Right? There's an empty space. We'd love to be here, but there's nothing there. So the best we can do is get close to that. And that's what I want to talk about next, which is this concept of non-dominance.

So what do we mean by non-dominance? What is a Pareto frontier and what is multi-objective optimization? The key point about this is that when you do non-dominance filtering, when you look for Pareto frontiers, you do not need to express your preferences ahead of time. Right?

So when you do you a weighted Pugh matrix or when you do a utility theory, you had to define ahead of time, before you did the scoring, what's more important-- fuel efficiency is twice as important to me as cost. Right? The weighting, the preferences had to be expressed *a priori*-- before you did the scoring. And there's arguments that that's not the best thing to do, but you can do it.

Here in non-dominance and multi-objective optimization, eventually you have to express your preferences, but you do it after you've scored the options. That's a really important distinction.

So let me give you a little history here on this. And what's really interesting is there's a Lausanne connection. OK? There's a connection with Lausanne. So the word Pareto frontier is

named after a an economist, who actually started as an engineer, Vilfredo Pareto. He was born in Paris in 1848 and then graduated from the University of Turin in 1870. Very much in engineering, civil engineering, his thesis title was The Fundamental Principles of Equilibrium and Solid Bodies. Something we take-- that's pretty basic now. Forces and torques in equilibrium on a body.

Now he worked as a civil engineer in Florence, and then he got interested in philosophy, politics, and economics, and he started to think about this in terms of how does this concept of equilibrium apply to economics? And in 1893, he actually became a professor at the University of Lausanne, which is located right next to EPFL, and started applying this to economics and societal theory.

And then he did-- some of his work was a bit controversial, but the one that I want to talk about here is this the idea of the Pareto optimum. And I will-- let me just read this quote to you. "The optimum allocation of the resources of a society is not attained, so long as it is possible to make at least one individual better off in his or her own estimation, while keeping others as well off as before in their own estimation."

So what it means is if you can make an investment or make a change in society that makes a particular individual or group of individuals better off in their own estimation, without having to take the resources from another group and make them less well-off, then you've not reached an optimal allocation of resources. You know, if there can be a win-win-- if everybody can be better off by making certain investments, you haven't yet found the optimal strategy.

Only when-- the only way to make somebody better off in their own estimation than somebody else is by taking-- and this is a big political debate still today-- by taking from one group and redistributing resources to the other. When you're in that situation, then you are at the Pareto optimal point. OK?

So that's the key idea underlying this. So what this means mathematically is that an optimal solution -- X^* -- is optimal if and only if-- for a feasible solution-- so, first of all the solution has to be feasible. And this is when you have multiple criteria. We're trying to do vector optimization, and x must be what we call an efficient solution, and it's efficient only if its objective vector-- J of X -- is non- dominated.

And what this means is that if you're looking at a point, it's only efficient when it is not possible to move from that point to another point without degrading at least one of the other objectives.

So if you can move from a particular design point to another design point, and all the objectives get better, you're not yet efficient. You're not yet at a Pareto optimal point. And this gets us to the notion of dominance.

So dominance-- let's say we have two designs, two alternatives-- J1 and J2. Those are their objective vectors. It means that-- this is for maximization. We're trying to maximize. It means J1 dominates J2 weakly if J1 is better or equal J2, and at least in one of the objectives, I, strictly better than J2. So you could have two alternatives. This gets back to Pugh matrix and the question-- who asked about the ties? You did, right? It's so long ago. It's like, wait a minute, it's like an hour ago, you know?

So we have two alternatives, OK? They're all equal. They're tied in all criteria, except for one, where J1 is better than J2. Which means that J1 will dominate J2, but weakly, because there's some ties. Right? And then there's a stronger definition, which is that J1 strongly dominates J2 if and only if it is better strictly better in all attributes. So for J1 to dominate J2 strongly, it's got to be better in all attributes than J2. Does that make sense?

OK. So I know this is a bit abstract. So let's do a concept question that's sort of based on a hypothetical but real example. So assume-- let's say you're in charge of designing in new commercial aircraft, and there's four criteria. You want to maximize range. You want to minimize the costs-- dollars per kilometer flown. You want to maximize the capacity of the airplane-- number of passengers. And you want to maximize cruise speed in kilometers per hour. So it's a multi-objective aircraft design.

You do all the work. You do concept generation. You come up with eight concepts for airplanes, and you evaluate them. And they're shown here-- one, two, three, et cetera, through eight. So let's just look at number one. What does this mean?

Airplane number one, concept number one, has a range of 7,587 kilometers. It has a operating cost of \$321 per kilometer. It can carry 112 passengers. And it has a cruise speed or max speed of 950 kilometers per hour.

So I'm going to leave this up. This is like a five minute-- this is-- on the next slide, I'm going to ask you-- and we'll come back to this-- which of these airplane designs, these eight, are non-dominated. Meaning that they're not dominated by any of the other. And we're going to apply weak dominance here.

That was going to be your question, right? Weak dominance. So I don't think it actually matters. I don't think there's too many ties in here, but apply weak dominance. So which of these are non-dominated? Meaning that you wouldn't discard them right off the bat, because there are some strong features that they have. So work through this. You might want to write down which ones are non-dominated. And then we'll do the concert question, and then we'll show you the answer.

AUDIENCE: I just want to make sure. So non-dominated as in there is not one that could be considered better than it?

PROFESSOR: Yeah. So, this is the definition. All right. So let me-- so please submit your answers. So here are the choices. Which airplane designs are non-dominated? Five, six, and seven. One, three, four, and eight. One, two, three, four, and eight. Two, three, five, and six. One, three, four, and seven. Or you need-- this is not answerable. You need more information to answer the question. OK?

So we have a pretty good distribution here. So the wisdom of crowds I think prevails. So the correct answer is the third one-- one, two, three, four, and eight. So 36% of you got it right. Now let me let me show you the solution to this. So it is, in fact, possible to answer this without additional information.

So the way to do this is you have to do pairwise comparisons. Now you said about-- I think your strategy-- What did you say, early on?

AUDIENCE: I was looking at maximums.

PROFESSOR: You were looking at maximums. So you were saying who's the best performer among the concepts for each criteria? And you found those pretty easily, I would assume, right? So what-

AUDIENCE: [INAUDIBLE]

PROFESSOR: OK. So you've got to get the sign right. Do you have the mic on, by the way. So what would-- if that concept is the best performer in a criterion, what does that tell you?

AUDIENCE: It means that none of the other performers have something equal to or greater than that performer for that criterion.

PROFESSOR: And this relates to the question-- is it Martin? Martin? In the red shirt, first row? I didn't get your name right. I'm sorry.

AUDIENCE: Bastian.

PROFESSOR: Bastian. Bastian. Bastian. You remember at the start of the lecture, you said-- what did you say at the start of the lecture?

AUDIENCE: I said that if you evaluate all the concepts based on just plus or minus for the criteria, you might have one where you evaluate all the-- a bunch of criteria that are negatives, but there's just one that is positive. But that one might be very, very positive.

PROFESSOR: Right. So that's exactly the point here. That's exactly this point. If you are-- if this concept is the best performer in just one of the criteria. Like it has the best speed, or the lowest cost, but it's terrible, terrible in everything else, it doesn't matter. In terms of dominance, it is non-dominated. It cannot be dominated by any other concept. If you are the best in class, even just for one criteria-- does that make sense? You cannot be dominated by any other design if you're the best performer on just one criterion.

So let's think about track and field. So the people-- who's doing track and field? Any athletes here? Like runners or javelin or-- Veronica. I didn't know that about you. Cross country and track. OK. So 5K? All right.

So let's say-- I'm sure you were very good at 5K. And how was your shot putt?

AUDIENCE: I never competed in the shot putt.

PROFESSOR: OK. So then we don't know, right? But the point I'm trying to make is you could be a super duper specialist in one criterion, and that's-- you can will win a gold medal with that. But you're very maybe not good in other things. So that's the point here is in order to be dominated, if you're the best performer-- so, right off the bat, in this exercise here, in our little exercise, you can remove all the concepts that are best in class from the set. Because they cannot be dominated. So they have to be non-dominated.

OK, so the way you do the scoring with pairwise comparisons is you say, let's compare one and two. Concept one and two. Well, one is better in range, right? And it's better in speed. But then the second concept is better in two and three. So the score here is two versus two, which means neither one nor two dominate each other, right? Even if it's-- what if it's three to one?

What if it was a three to one score? It would still be true, right? That one doesn't-- one is, in a sense, better because it's better than the other. But it doesn't dominate it, because dominance is a very crisp definition. Now this is different, right? Compare one versus number six. Concept one versus concept six.

One is better in all four criteria to six. So clearly, solution one dominates solution six, and as a rational decision maker, you can eliminate concept six from the set because it is completely dominated by concept one. OK? So in order to be dominated, a solution must have essentially a score of zero in a pairwise comparison.

Now, if we apply this to the full set of eight we can actually show this as a matrix. I call that the domination matrix. It shows which solutions dominate which other solution-- the horizontal rows and the vertical rows. So the way you read this is so where you see these dots, these are dominance relationships in the pairwise comparison. So this tells you that solution two dominates solution five. This tells us, if we look column-wise, solution seven is dominated by solution two and solution eight.

So, if we do the row sum-- if we sum this along the rows, the row indicates how many solutions this particular solution dominates. And-- but the question was about non-dominance, so we need to look at the column sum. The column sum indicates by how many other solutions the kth solution or concept is dominated. And when you do this column sum, you can see that concepts one, two, three, and four have a zero. They're not dominated by any other design. Column eight has a zero, and concepts five, six, and seven are dominated each by at least one other concept, which means that they're dominated, and the others are non dominated.

So when you do your concept selection, you can actually apply this to a much bigger set. You filter out all the dominated solutions. And it turns out the bigger your number of alternatives or architectures, the smaller percentage-wise the set will be of non-dominated solutions. And those are the ones you really want to focus on. Is that clear? I know this takes a little thinking about, what does this really mean? But but that's a very rigorous way to do it.

And then this gets us to, eventually, this gets us to the notion of Pareto optimality. So what's the relationship between Pareto optimality and non-dominance. So let's say we have two objectives, J_1 -- we want to maximize J_1 . Maybe that's performance, endurance. We want to minimize J_2 -- maybe cost, for example. So the Utopian point, or the optimal corner to be in, is the lower right. We want to maximize J_1 , minimize J_2 . There's nothing here. It's empty.

There's nothing feasible here.

So the best we can do is get close to it. So when you look at a set of discrete points, discrete concepts, those concepts can have three properties. They can be D, meaning dominated. So all the red points shown here are dominated. And the way you can see this if you draw a little-- let me try to let me try to draw a box here.

Where's the eraser?

So let's pick a point, a particular point-- Let's pick this point right here. Why do we know that this point is dominated, just looking at it graphically? How do you know that? Go ahead.

AUDIENCE: The point is not at the front.

PROFESSOR: It's not at the front, but the way the way to find first of all, whether it's dominated, which is relative to the other solutions, is we can just draw a horizontal line here, and a vertical line here. And so in a multi-dimensional space, you're essentially drawing a hyper-cube. And down here is our utopia. This is where we would love to be, but we can't be. It's not feasible.

And you see in this box, between the point and the utopia, there is this other guy here. This point here. And because it's there, it's closer to the utopia, and it dominates this point in both objectives. Let's take this point here. Let's do the same thing. Draw a vertical line and a horizontal line. And there's a whole bunch of points, all these points right here are all the box, closer to the Utopian point. So this point here is actually dominated by one, two, three, four, five, six seven other points.

So that's essentially dominated. Let me erase this again. Then we have these darker red, purple points. Let me just circle them real quick. So we have this point here, this point, this point, this point, this one, and this one. And those are what we call non-dominated, like we just did the filtering, meaning there there's no other point in the set that's better than it in all criteria. So you would consider them as concepts to choose from.

But they're not exactly on that dashed line, which means that there is a way to improve them. There's a way to tweak them, to optimize them further, if you could to get them right onto that dashed line. They're close but they're not exactly on that dashed line. So you could tweak these concepts to get them closer to the Pareto front. But they're still pretty good, because they're non-dominated.

And then we have those grey points, which are shown here. There's three of them. Let me use a different color for these. Let me use a different color, maybe green. Here. This point, this point, and this point, they are Pareto optimal. Meaning they're on the Pareto front. There's no way to improve them, such that you could make both objectives better.

Right? If you want to make J1 better, you have to sacrifice on J2. So the only way to modify these points is to slide along the Pareto front, meaning you're trading off the objectives. Do you see that? And so, in my spring class that I teach, called Multi-disciplinary System Optimization, we get into this pretty deeply. Like how do you actually do this?

But the point I want you to keep for now is that when you select the concept, do not select a dominated one. Get rid of the dominated ones. Filter them out. And then you still have to choose among the others, and that's where your preferences come in. Right, so given the blue and green choices here, which one are you going to recommend at the end? And that this is sort of the preferences between maximizing J1 and minimizing J2. That's a social process of coming up with your preferences.

But you now determine those preferences after you've scored all the concepts. Do you see the difference? That's pretty fundamental. OK. So, one more thing and then we'll wrap it up.

So now you've gone through all this work. You've done your concept generation, concept selections, scoring, and you found the one concept that you're going to go forward with. That's when you do PDR.

So what is the PDR? The Preliminary Design Review. You explain what concept and system architecture was chosen. You explain why it was chosen, and you want this to be approved as the design baseline for going forward and doing more detailed design work. You compare it against the rejected alternatives. Typically at the PDR, you will spend some time on saying, these are the other interesting concepts that almost were selected, but we didn't because of these reasons. You show some quantitative analysis that gives confidence, especially you want to show that the requirements that you had defined and agreed upon at the SRR can be satisfied by your chosen concept.

There's maybe still some risk or some uncertainty. So any risk reduction experiments or prototypes. So you can actually do prototyping at this stage. You don't want to spend too much on that, but you want to do enough of it so you can reduce risks. And then you do a preview of

the detailed design phase between PDR leading up to CDR. So here-- I'm not going to read this-- but this is the description of what the PDR is all about. In the NASA system engineering handbook on page 177, there's a very detailed description of entrance criteria, exit criteria, et cetera, et cetera.

So here's a description from a PDR from a European program, and you can see it's a pretty big group, right? This is like 30, 40 people. And because the PDR is such an important decision, that's when you also invite external stakeholders. And you want to make sure that people come away from the PDR with confidence that, yeah, we have a good system architecture. We've picked the good concept, and it's worthwhile now putting all this work into getting it to a detailed design that can actually be built.

OK. Any questions about PDR, or any experiences from PDRs that people want to share?

AUDIENCE: I have a question about the PDR. So you said there's 30 people in that picture. How do you reconcile that with seven plus, minus two rule?

PROFESSOR: Well, it's seven, plus or minus two, squared, right? So you have a hierarchy of people. Not everybody is a decision maker here.

AUDIENCE: I see.

PROFESSOR: So you probably have people who did some of the thermal design, some people did the cost modeling. So, it's a great question actually. You want many, many voices and people who contributed to make sure the PDR presents all the information you have. You don't want to do a PDR with just seven people.

AUDIENCE: But I guess you could say that in the presentation, I guess, only seven plus or minus two, will people be driving at that point in time.

PROFESSOR: So there will be a leadership, right? There is an organizational issue. But the point here is that it's a big review. It's a big milestone. Please, go ahead.

AUDIENCE: The big issue here is that all the people who are not decision makers, that are in the picture, they will have to implement the detailed design, so it's very important to have them in the process. Otherwise, you have to go back to your organization and do the PDR again to explain the what's and the why's. This is why you try to have these very large meetings. When you do a pre-PDR, you do a dry-run, and you send stuff in advance and make sure that you're really

all agreeing finally, if possible, you get the discrepancies off the table. But then you all agree that this is how to go forward and then the worker bees can go and do it.

PROFESSOR: Yes, that's another reason for being inclusive. I agree with that. Any other-- who has gone through a PDR in their prior work? Anybody want to share? Marissa. Tell us about how that went. I mean as much-- I know you can't probably we share everything but--

AUDIENCE: For us, I mean, it was a pretty big milestone. You work a lot leading up to a PDR and then you finally have all the really big stakeholders in the room. And a lot of them have been involved during the design phase, but when you actually have them all sitting there. And I think getting the buy-in was also really-- I think that's a really big part of PDR. Like normally you're feeling pretty confident in your designs, but then actually making sure whoever the agency is that you're working with, then your stakeholders are also kind of all on the same page. It's pretty critical.

PROFESSOR: Did you discover problems or discrepancies during the PDR?

AUDIENCE: Yeah. I definitely think-- just generally we've had delta PDRs, if things weren't completely thought through. Or you'd have a fair number of action items that come out of PDR to go and address before you really are allowed to start moving towards CDR.

PROFESSOR: Good. Any comments or experiences with PDRs at EPFL?

AUDIENCE: I did a PDR at the European Space Agency during the REXUS BEXUS program three years ago, and as far as I remember it was very appreciated to bring as many details as possible. And to be very close to that critical design review stage in the PDR.

PROFESSOR: That's interesting. So, that's good. I think there's a good to have detail and so forth. But what's the downside? Can you think of a negative what would be what would be not so good about almost being a CDR right there? Can you think about a downside to it?

AUDIENCE: Probably not taking too much time in order to clarify the concept and the requirement because you are in a rush to finish the product.

PROFESSOR: That's a good point. So it takes more time. But what I'm thinking about is it's very difficult to then undo the decision or backtrack. Let's say that at the PDR that there's some-- it could be that you really forgot something very important. And that your concept is actually flawed. But you didn't know it as you were working inside your team. And then if you have a design that's

almost CDR quality, it's very difficult to undo that. So I would say that a PDR-- I actually believe that at a PDR you should have some detail, but not too much yet.

I think a true PDR should be at the system architectural level, with some detail but not all the details. But thank you. That's a great comment. Thank you very much.

So let me close it here. So in terms of concept selection during conceptual design, we can use Pugh matrix selection. And that's essentially useful when we don't yet have those detailed mathematical models. It's qualitative. And then as we move more into preliminary design, we use utility analysis, especially for noncommercial systems where a net present value-- a financial-- in commercial companies, you will often do a financial analysis as a measure of utility.

And then this concept of non-dominance is very, very important. Generating many designs and then applying Pareto filters to try to come down to the non-dominated set. So that's essentially the essence of it. So I hope that was clear. A-4 will be posted later today.