



Make your choice!
**- On the empirical
exploration of good
decisions**



**English
podcast transcript**

Everything is nothing

An introductory podcast to Daoism

English transcript

Zürich: METIS Podcast Transcriptions 2023

Translation by Eliane Schmid

ES = Eliane Schmid

RH = Ralph Hertwig

ES: Hello and welcome to Wisdom Talks, the podcast accompanying the METIS project, the internet portal for intercultural wisdom literature and wisdom practices to be found on www.metis.ethz.ch. My name is Eliane Schmid, and I will be hosting today's podcast centered on empirical wisdom research.

During our Wisdom Talks, we often talk about cultural conventions, famous authors, and books from various wisdom traditions, but there's one perspective that we have not yet discussed explicitly, and that is the perspective of science. Wisdom traditions often aid us in self inquiry, reflecting on ourselves, our lives, or on a good life. They often position themselves in relation to prevailing knowledge claims and worldviews. It is not uncommon to talk about experiences that affect every human being at one time or another, but are wisdom traditions actually empirical endeavors? And how can wisdom or topics of wisdom, such as death, aging, decision-making, et cetera, be the subject of empirical research?

I'm very happy to welcome Professor Dr. Ralph Hertwig at our ETH Zurich Studio today. He's the Director of the Center for Adaptive Rationality at the Max Planck Institute for Human Development in Berlin. Thank you very much for being here today, Ralph, and for discussing these interesting but also difficult questions on research on empirical wisdom.

RH: Well, thank you very much for having me, Eliane. Hello, and I'm looking forward to our conversation.

ES: So then let's start directly because I'm very curious to learn more about this. Since Solomon actually, the ability to decide correctly, if we can say it that way, has been considered an indicator of wisdom. Today we may speak less of wise decisions and more of good, or reasonable, or well-reasoned decisions. What constitutes rational decisions has now become the subject of empirical sciences such as sociology, behavioral economics, and psychology. These disciplines use experimental methods to investigate why people choose to do something in one way or another. Could you perhaps briefly tell us which theories were and are particularly significant in this area, and how does one explain human decision-making?

RH: Sorry for laughing, but this is such a simple question. I don't even know where to start. Maybe let's start with the notion of even making correct decisions because that's by itself a very complicated notion. Because in some cases there is a ground truth and you know exactly whether something is right or wrong simply because there are facts in the world against which you can measure your judgment. But then of course there are many preferential judgements, a matter of taste so to say, where it's very hard to say what a correct decision is. And last but not least, there are of course also ethical dilemmas that we encounter all the time. If you are thinking in the context of the pandemic, but also far beyond that, where it's also extremely difficult to say what is right or what is wrong, and depending on your perspectives, you may arrive at very different judgment, decisions, or choices.

Now, to go one step further, I would agree with you that wisdom per se is not an enormously well studied topic within psychology. Rather, what is much more studied is what is a good decision, what is a reasonable decision, or very often what is a rational decision. And in terms of rationality, there are some prime theories that are being considered as the theories that codify good decision-making. And they often come from economics, such as, for instance, expected utility theory, or subjective expected utility theory, game theory would be another example. And also issues such as probability theory, or statistics, and the norms that can be borrowed from that, such as for instance, probability updating in a Bayesian way would also be considered to be a benchmark of good reasoning. That's just a first attempt to answer your question.

ES: I would like to dig a bit deeper there and question the idea of a good decision. Is there a scale of good decision-making, or is it contested kind of way of defining this and it's better, as you said before, to talk about rational decisions?

RH: To the best of my knowledge, there isn't a scale, but what very often is done is that decision situations are being created where there is a normatively correct answer, where for instance, it's a

choice situation, you make a choice between two alternatives, and to the extent that you are a rational expected utility maximizer, or that you are a rational person in terms of game theoretic concerns, we know exactly what the correct answer would be or not. And then we can basically observe – and that's the traditional way of going forward here – we can observe to what extent your answer conforms with such a norm. Sometimes it's a binary judgment, it conforms or doesn't, and sometimes there are judgements, continuous judgments where we can quantify the extent to which you arrive at the normative answer or not. So in that sense, we can scale it in some cases, but very often, it's just you conform to the norm or you don't.

ES: So it's also very much dependent on the field of study you're in, or the topic. You mentioned economics for example, but then in everyday life you would have different scales, if I use this term again, to judge if something is a good or a bad decision.

RH: Yeah, I mean that's an interesting question because if you come from a normative theory such as expected utility theory, they wouldn't make a distinction between the gambling choice situations that are often used to study rational decision-making and the important decisions that you make in real life, such as whether or not you have children for instance, or whether you get married, whether you buy a house, what career choices you take. And here the view is that these normative frameworks in principle are so general and so powerful that they can apply both to very minute situations that we often study in a laboratory, but also to the big situations in real life. So the assumption, the theory is applicable to small and big decisions that we make.

ES: I'm going back a bit also to the very beginning because you are the Director of the Center for Adaptive Rationality. What is adaptive rationality?

RH: That's a very good question, and that brings us back to a theory that is domain-general such as expected utility theory or game theory. And as I explained before, the assumption here is that the theory itself can apply to all kinds of situations, and in this way it's a very powerful but also to some extent “context-inspecific” theory because it doesn't care really about the context. And my argument and argument of many other scholars in my field has been that these theories are too general in order to really capture reasonable behavior. And let me give you an example. So expected utility theory or game theory is built on the notion of coherence. That means that in our decision-making, we should be, for instance, internally consistent. If we are internally consistent, then we are reasonable, then we are

rational. But that overlooks often other important concerns that are at least as reasonable as the maximization of expected utility theory is.

And let me give you just an example to make this a little bit more tangible. So just imagine that you are at a party and you are sitting at a table and at the very end of the dinner, basically a fruit plate is going around the table, and finally it arrives at you. And there's one person after you, and what you see, there's only one apple left on the plate. And you, being a very polite person, as I can tell, you would say: "Well, I'm not going to eat the apple. In the choice between choosing nothing or the apple, I prefer nothing, and pass the food platter on to the next person."

But now the host suddenly sees that, oh my god, there's only one apple left: "I'm adding another apple to your platter." And now again it comes to you and you have now the choice between two apples and choosing nothing, and suddenly your preference changes because now you are choosing the apple over nothing. From the perspective of a coherence based or domain general theory such as expected utility theory, the normative principle here is called 'property alpha'. I would say that wow, you are irrational, but of course you are not because in one case, your politeness tells you: "Well, I shouldn't take the apple" and in the other case it was clear that taking the apple is consistent also with your concern for politeness. And in this sense, and that's what we mean with adaptive rationality, that from situation to situation, from context to context, your concerns, the benchmarks of rationality may change, and therefore these context unspecific domain general theories are often very problematic when we used them in order to paint a person as irrational.

ES: But just in this specific situation, you said that I would not take the apple because I'm polite, but this is very much also a cultural thing and also a character thing in the sense that maybe another person is polite, but they're like: "Well, I want an apple, I will take it," but that does not necessarily make them impolite. But depending on the cultural context you're in, some people would say: "Oof, like she's a woman and she should pass on the apple to the man sitting next to her because he's also older, and she just took the apple for herself, how rude is that?" In another context, it would be like: "Well of course you should take the apple, she's hungry and she needs that apple." How would you bring in cultural context, and personal characters, and other similar layers to this?

RH: And as you described, your behavior may be at first sight inconsistent across different situations, but if you look a little bit closely, we may be able to explain it. Like for instance, in one case you are really hungry, and hunger in this case really beats your politeness. In the other case, you are not hungry and

you can easily say that now the norms of politeness are more important. And it's just another example that people are highly flexible and they can switch back and forth between different criteria and different benchmarks for their behavior, which is another example of adaptive rationality.

But coming back to your question, you are of course totally right, in particular if you think about social norms, but also for instance about self-concepts, which is also important. My self-concept also informs my behavior and informs whether my behavior is reasonable or not, or whether I consider my behavior reasonable or not. And these are all concerns, self-concept, cultural norms, social norms, that are not captured in these dimensional theories of rationality such as, say, expected utility theory. And that's why we think that we need to make these frameworks richer. We need to go beyond these frameworks. We need to think about adaptive rationality.

Another way to answer your question is they are indeed studies that try to understand good decision-making across different cultural contexts. Not too many, but there are some. There's one complaint about a lot of studies in psychology but also in economics, that a lot of the studies are conducted with the so-called WEIRD population, WEIRD stands for Western, educated, industrialized, rich and democratic. And it's true that if you look at the samples in the population of many of the behavioral studies that are being conducted to figure out whether people make good decisions or not are often run with these WEIRD populations.

But there are also some studies that try to look into small scale societies, which is of course much harder to study because you actually have to go in the region where small scale societies still exist and live. And then it's interesting to see what kind of results are being produced. So for instance, there are studies that try to understand the norm of fairness, for instance, at what point does fairness, which is also a concern for reasonable behavior, when does this normal fairness evolve? And one of the insights of these small-scale society studies is that the moment a society or a small group of people exhibits something that has been called market integration, which means an exchange of goods entering almost commercial relationships with other people, then the normal fairness becomes actually pretty important. And one of the speculation is that has something to do that in market exchanges, your reputation is particularly important for the success, in particular, for repeated exchanges.

So, this as you can see, looking at culture, looking also at society and societal institutions is very important to understand what good or bad, or reasonable or unreasonable behavior is. What I would say is that we must not only look at the individual, but we also need to look at the individual in the cultural, and social, and institutional context.

ES: I'm quite amused, I must say, it by the concept of the WEIRD, and maybe it is fitting, maybe not, but I'm now curious, are there some decision-making principles that you can see across cultures that every culture agrees on is a reasonable way to make a decision? Is there something that combines multiple cultures that they're all on the same level, or do you think that it depends on multiple factors such as age, heritage, again, culture, state of knowledge?

RH: Again, one of these very modest questions that require almost a lecture to think about and to talk about it. But let me try to give one possible answer. And if you ask me tomorrow, I would probably give a slightly different answer because it's such a rich question that you are asking. Of course, all the factors that you mentioned play a role in decision-making in the actual processes. For instance, people make different decisions across the lifespan. That has a lot to do with the fact that our cognitive faculties, that our cognitive resources changed both from young age to an adult, but also of course when we get older and our cognitive capacities tend to decline. And this has implications for the way we make decisions also for the speed that we make decisions, and we already talked about the cultural embeddedness of decision-making. So all these factors, and many more, make a difference.

Having said that, there are common structural features of decisions that cut across cultures, that cut across situations. So to give you one example, many of the decisions that we take, we take under conditions of uncertainty. Everything that concerns the future is to a large extent unknown or uncertain. Similarly, a lot of the decisions that we take, we need to take under time pressure. We don't have all the time in the world to research all the options that there are. Let's just say that you go to the supermarket and think about how you make your decisions. You don't have all the time to look at all the options to think about which one fits your preferences best and satisfies your needs the best. But you make decisions often under limited time and time pressure. And what we do know from a lot of research also across cultures is that time pressure changes the way we make decisions, often then we make it in a so-called heuristic way, that is, we rely on very little information in order to make decisions.

So given that a lot of the situations in the world share these structural features such as time pressure, such as limited amount of information, such as uncertainty, we can safely assume that that there will be commonalities in the way people across different cultures make their decisions, simply because the world imposes on us a certain style of decision-making.

ES: Are rash decisions or fast decisions necessarily worse decisions than if we have a lot of time to think about it, or could it at some point also change and thinking about something for too long a time will make it an irrational decision as well?

RH: Yeah, and that's actually a very important question because that basically is a bone of contention for, I would say, almost the last 50 years in a lot of research in psychology. And there was a program that started in the early 1970s that was called the "Heuristics and Biases" program. It's associated with two very famous psychologists, Amos Tversky and Danny Kahneman. Danny Kahneman is very well known because he wrote a very famous popular book, *Thinking Fast and Slow*. And the idea in that research program was indeed that when people make heuristic decisions, which is often associated with decisions that are based on relatively little information that are done very fast, or fast at least, and where the processing of the information is not complex, that such decisions can easily lead to systematic errors and biases. And in this research tradition, heuristic decision-making or what you called rash decision-making is associated with not so good, often irrational decisions.

That's one view, but a more novel recent view, but even that view is now already 20 years old, and that was researcher that was started by Gerd Gigerenzer, who actually was my predecessor as one of the directors at the Max Planck Institute for Human Development, and we are building on that program, on so-called fast and frugal heuristics. That research program that Gerd Gigerenzer started and we are continuing, has shown over and over again that very often relying on one or two pieces of information is actually a very powerful way of making decisions, and often can beat a much more complex, and statistically or otherwise more complex decisions.

And let me give you one example, and that's funny because I just read about this paper this morning. So one of the issues that is very important to us but is actually an extremely difficult decision is to figure out when somebody's lying at us. And you can imagine with all the talk about false information and with deep fakes and so on, the issue of what is right and wrong becomes increasingly important. But as we know, training, for instance, experts, or training algorithms or computers to figure out when somebody's lying at us is extremely difficult. People are not very good. And one of the reasons is that a lot of information where we believe that could potentially be predictive of somebody lying is actually not good, is not valid information. So for instance, if I look into your eyes, then clearly I'm not lying at you. Well, it's not a good cue if you use that cue. It turns out, and that is based on empirical research, that there is one relatively valid and powerful criterion and that is the detail of the narrative that the person who tells me something is choosing.

And it turns out that people who tend to lie use less detail in the narrative. People who tell the truth are using much more detail. And that is actually, and that has been also shown empirically, a very simple cue that can quickly be evaluated by people. I can quickly pick it up. And what has been shown in this paper that I read this morning is that if people are taught this one cue, only that one cue, they are much better than other people who use all kinds of cues, many of them are wrong, and try to combine them. So this is just one example that heuristic decision making can be really very efficient and good.

ES: There's a fine line here though to what we would also call gut feeling or people often give you the advice to follow your heart. Then when we talk about lying, people often have these gut feelings that there's something not quite right about this story or I feel uncomfortable. And we would call these maybe irrational decision-making tools, but where do these feelings come in? Do you analyze that as well, or not necessarily?

RH: Yeah, intuition and gut feeling, these are these concepts that are a bit vague, and we try to understand them better. And for instance, one way to understand intuition and gut feelings is in terms of heuristics, because very often we are not necessarily aware that we are using specific heuristics. In the example that I told you, people were explicitly taught a heuristic, but many heuristics that we are using, we couldn't necessarily spell them out, and yet we are still using them, but we call them gut feeling, or we call them something like intuition. Now an important question, and that's the one that I also hear in your question is, under what circumstances is reliance on gut feelings and reliance on intuition a good thing, or is likely to lead to good decisions, rational decision, or even potentially wise decisions? And when does it lead astray?

And there are a number of constraining factors, but let me just tell you briefly about one, and that is the extent to which you were able to collect and to muster experience in the environment in which you are using your gut feeling. So for instance, to give you one example, just think about personnel selection, and you are the person who is responsible of picking a person for the job in your company, and you decide in favor of some people and you turn away other people. Now, this is an environment in which you will never get feedback about the person whom you turned away. They just disappear. You will never find out whether your intuition that they are not as good as the others was good or not. And what that means is that sometimes the environment only gives us partial feedback, and then it's much harder to form a really good intuition, a good gut feeling.

Or think, for instance, of a doctor in a, say, emergency unit who immediately tries to treat a patient and then this patient is being moved to the next ward, and you will never see that patient again. And so you may never learn whether your early diagnosis and even your treatment was a good one or not. So an institution can change that, an institution can make sure that you get all the feedback in some cases, of course not in all cases, can make sure that you get the feedback in order to then really evolve and form good intuitions and gut feelings.

ES: Now, we talked before also about this fast and slow decision-making, you talked now about that we don't always have feedback about our gut feelings. Would now be interested in the after, in the sense that, are people happier when they have the feeling that they had enough time to make a decision, or not necessarily? Can they find more peace with a bad decision when they think, "Well, I handled under stress, I did the best I could, and this is the outcome." Is there a relation here, or not really?

RH: I'm not aware of any research that looked at that, but I would assume, that's just a speculation, that of course the extent to which we are happy with our decision has something to do with the outcome of our decision. If the outcome is desirable and is consistent with our values and our goals, then of course by definition we tend to be happy with it. But sometimes the decision doesn't turn out the way we want it, and then it is important that we are happy with the procedure that we have chosen to make the decision. If we have the feeling like: "Oh my God, I just should have thought a bit more about it," or, "Oh my God, I should have paid attention to something else," then of course we also may be dissatisfied because the process wasn't the one that we wanted to live up to.

So I think that there are various dimensions that we use in order to then inform the level of happiness and satisfaction with our decision. But of course, we also fully realize that sometimes decisions are done under enormous time pressure. And then maybe that's another important concern here is that often we are not aiming for optimal decisions or the best possible decision there is, but we are aiming for good enough decision. That's actually a very important concept also for our individual happiness and wellbeing. Because just imagine in many areas of our life, there are so many options that if you wanted to always make the best decision, we would never get to make a decision because we would always be stuck in finding out what is the best option for us. So people who try to optimize, who try to get the very best thing for them are often the people who are less happy with their decisions. But people who are content with good enough outcomes who say: "Well, it's maybe not the best outcome, but it's good enough for me," they tend to be happier with their decisions.

ES: Now, you talked a lot about the individual, about good enough for me, and making a decision for myself. What about, now I'm also going in the direction of social media, for example, we're actually in a big network, the decisions that we make often don't only have an outcome for ourselves but influence a lot of people. How do you combine the individual and the crowd? How do you bring the decision-making concepts for yourself but also for a big group of people together?

RH: That's another one of these big questions. I think you have to be more specific in your question, what's on your mind? What would you like to know?

ES: So for example, we try to make decisions for ourselves that are good, but then we also see guidance, for example, when we look at Twitter or Instagram, because we want to know what others are doing and conform to the standards that are set. This is in the direction that I want to go. How do these big media platforms inform our decisions and how do we make decisions in accordance with these platforms to fulfill some kind of norm?

RH: Now I'm getting the drift of your question. So first of all, let me say, going back again to the notion of good decisions, that very often in particular when we deal with uncertainty, which is an ubiquitous condition of our decision-making, that one way of coping with uncertainty is not necessarily to try to figure out everything by yourself, but really benefit also from, if you will, the collective wisdom and the information that is distributed in your close environment. To just give you a very simple example, when you go to the restaurant that you've never been to before and you do wonder: "What the hell is a good choice of the dish tonight here?" Well, let me ask you, how would you make that decision? Would you sit there and look at the menu, study it for minutes, figure out all the options, thinking about whether it meets my preferences or not, and after 15 minutes you finally made a choice. Is that how you would do it?

ES: I'm often very hungry and go with the gut feeling, as I mentioned before, and just pick the thing that will feed my belly in the fastest, most efficient way.

RH: You are a real connoisseur, I can tell. Here's a different way of doing it, and this is if you will, a socially rational way. So one thing you could do is you could turn to the waiter and ask him or her, not what he or she would recommend, that's the wrong question, but ask him or her: "What would you eat here tonight?" Because that basically circumvents a lot of other concerns that the waiter or waitress may have. Or alternatively, what you could also do, you could look around and see what most of the

people actually have chosen on their plates. And what I try to illustrate with that is that often a social environment is a source of at least information, sometimes maybe even wisdom, and it's really smart to use your social environment.

But coming back to your questions though, in principle, using your social environment to inform your decision-making is a very smart thing to do. But now, as you say, we are now often navigating digital environments where the social environment is a very different one than the one that we encounter from face-to-face. And what that means is that sometimes we need to be much more critical with regard to the social information that we are picking up in the digital world. Because for instance, it's not just the preferences of other people that we're seeing, or the opinions of other people, but there is also algorithms at work because what we see is filtered through the operation of algorithms. And of course the algorithms make rather select information according to totally different criteria, not according to the criterion of information utility, but for instance, according to the criterion of what is most likely to capture my attention. So in that sense, trying to learn from that social environment in a digital world, it's basically a different kind of challenge than in the world outside of the digital world.

ES: And now we've breached the topic of digital, and now I would like to go deeper in the topic of artificial intelligence. Now, artificial intelligence like ChatGPT subsumes a lot of things that are found on the internet, a lot of decisions that were made, a lot of information that are jumbled together, how can artificial intelligence in a way be a tool for us to make rational decisions? Can we just outsource rational decision-making now and think like: "Oh, they subsumed all the proper knowledge in the world and I can rely on that tool to help me?"

RH: In reality, decisions and also important decisions are already made by algorithms. There are many contexts in which we use algorithms to make decisions, and in some cases this is really efficient and really good. For instance, algorithms are really good in terms of pattern detection. So for instance, if there's a mammogram and you need to figure out whether the lesion that you are seeing here is one that is benign or whether it may be indicative of cancer, then it turns out that pattern matching algorithms or pattern detection algorithms are really good, and maybe even better than expert doctors who make the decision. So there are contexts in which we already made a lot of progress, and it's increasing evidence that algorithms can make very good decisions.

There are also other contexts in which we allow algorithms to make choices such as, for instance, allocation of resources according to certain criteria, so for instance, welfare benefits that are being

allocated according to certain criteria. And here of course, we need to be totally aware of certain traps that we may be walking into, because the algorithm will make this decision on the basis of input data that the algorithm is being provided with. And that input data may again be, of course, full of the same kind of biases, stereotypes, or other pieces of information that bias our decision-making when we do it on a personal level, and what that means is that the algorithm may be reproducing the same kind of biases, for instance, in the allocation of resources than the individual person would do. And what that means is that only because an algorithm does the decision doesn't necessarily mean that is a good decision or it's a decision that is free of biases. And that is an enormously important topic also for us psychologists to study. And there we are also now beginning to think about how we can, for instance, blind algorithms for certain kinds of information that in all likelihood will bias the algorithm.

So the answer is, there are domains in which clearly algorithms can make good decisions, but there are also domains in which we have to be very careful and think hard about whether the algorithm makes good decisions.

ES: Now, a good decision or a good enough decision on my part is now to tell you that time is running away. But I have a last question that I'm very interested in. Do you have a tip for us, for the listeners, just an easy tool, maybe that's, again, not an easy question, but do you have some kind of tip on how we can make a good decision in our everyday life, when I'm choosing between an apple and a peach?

RH: You are really asking the simple questions today. I think from what I told you, you probably figure that my answer would be, it depends. It depends on the situation, it depends on what your goals are, it depends on whether you make fast decisions or not. In that sense, there is no one golden way of making a good decision. And maybe my tip is to really think about in this situation, what would really help you to make a good decision? And in some cases it is talking to your friend, or talking to your parent, or talking to an expert. And another case is throwing a die when you're choosing between two options, and when you feel when the die is coming down, or the coin is coming down that you already feel like you want one option to be preferred over the other, then you already know your choice. So sometimes it's a chance device, sometimes it's the advice of an expert, and sometimes you have to really research, for instance, the internet or other sources and really think about it. Sometimes it takes time to make good decisions and sometimes it doesn't.

ES: I think these are very nice words to end this talk, and I would like to thank you for coping with my not-so-simple questions, but I learned a lot and I hope our listeners did too. So, thank you very much, Ralph.

RH: Thank you very much for having me.

ES: At this point, I would also like to invite our listeners to follow further Wisdom Talks as well as to curiously plunge into the multitude of texts and further podcasts that can be found on our website, www.metis.ethz.ch, the internet portal for intercultural wisdom literature, and wisdom practices. You can also find more information in the show notes.

Thank you very much for listening and goodbye.

This METIS Wisdom Talk was produced by Martin Münnich and supported by ETH Zürich and the Udo Keller Stiftung Forum Humanum in Hamburg.