Uncertainty for uncertain decision makers

PhD Thesis

Malvina Ongaro

Università del Piemonte Orientale
FINO PhD Programme
Academic Year: 2021/2022

Supervisor: Prof. Vincenzo Crupi

Contents

Introduction

Even though I usually say that it was good teachers and intellectual curiosity towards something new, it is possible that what really brought me to the philosophical study of decisions was my chronic indecisiveness. Whether it is the choice of a new flat or the selection of an ice cream flavour, I can spend disproportionate amounts of time pondering my options, weighing pros and cons, and preemptively regretting a choice I have not even made yet. In decision theory I found a clear account of what a decision consists in and what its variables are; and consequently I also learned that I could represent my uncertainty and its severity with probabilistic language, which meant that once I knew how to form and treat those probabilities rationally I would also rationally solve my indecision. In all its elegance, that felt like an account of decisions for very decisive people, one that would not really help me make all those big and small choices with which I struggled so much. Most of the times, my indecision is not (at least, not entirely) due to ignorance of some relevant contingency or future possibility: there is something else to my uncertainty. There is uncertainty about what I want, what is at stake, what I should prioritise, which options I have – and probably others, depending on context.

What is more, all these reflections happened at very uncertain times. Between political instabilities, the looming effects of climate change, and the insurgence of a viral pandemic, we face great challenges as communities, societies, and species. During uncertain times, the importance of experts in our collective decision making increases, as we want to learn as much as possible about our difficult circumstances in order to implement effective measures to face whatever is going on. The less we understand, the more we turn to experts in search for answers; and yet, under severe uncertainty there are many answers that the experts cannot give us, because there are limits to what even they know. Our need for trustworthy information may then hide the risk of over-relying on experts. The challenges we face are complex issues, the resolution of which requires both empirical knowledge and value judgements. If our uncertainty covers both, but we rely on scientists to solve it, then we are leaving to the scientists to make the value judgements too. Then, over-relying on experts can lead to shielding policy makers from their responsibilities, and to nourishing anti-science positions that perceive the political content of scientists' opinions.

Thus, it is for both personal and political reasons that I found myself investigating the concept of uncertainty, in its variety and in its role in decision making. This work is an attempt to understand what sort of uncertainty decision makers face, and to explore the implications of this analysis for private and public decisions: my main thesis is that there is a plurality of types of uncertainty, not all of which can be reduced with empirical methods.

The first chapter is devoted to a theoretical account of uncertainty and of its variability. In mainstream decision theory, uncertainty is expected to vary only in terms of severity, and the standard probabilistic treatment has been expanded accordingly to tackle increasingly severe types of uncertainty. But this approach assumes that uncertainty only concerns how things are, whereas there is a variety of doubts that agents can face in their decisions and that go beyond that. I propose that we understand uncertainty as a matter of conflicting reasons to hold alternative attitudes, which means that uncertainty arises when there is some disagreement and that the resolution of uncertainty requires the resolution of this disagreement. But given that in some cases this disagreement can persist even under ideal conditions, then in some cases the improvement of these conditions – for instance, with the collection of new evidence – cannot be expected to lead to the resolution of the uncertainty. On this account I construct a typology of uncertainty, identifying the conditions under which uncertainty is grounded in radical disagreement and retracing the role that these radical uncertainties can play in decisions.

Among the uncertainties at play in decisions, one which has received relatively little attention is uncertainty over how the decision is modelled. Typically, decision theories require agents to model their situation identifying a set of alternative options, a set of possible states, and a set of consequences. But what should be included in each of these sets is neither uncontroversial nor inconsequential, given that different models of the same decision may lead to different results. The agent may leave something out because they are unaware of its possibility, but also because they do not deem it to be relevant to their decision. While the concept of unawareness has started to attract some research, the notion of relevance is still poorly understood. My second chapter is devoted to clarifying its importance in decisions and to propose a possible interpretation of its meaning.

The last two chapters try to apply the analyses developed in the first two to policy making. In chapter three I use some of the types of uncertainty proposed to identify the boundaries of expertise in policy decisions. Understanding these boundaries is critical to provide an assessment of the "Evidence Based Policy" movement, an important trend in development economics that aims to increase the quality of policy making by focusing not on ideologies but on the testable effectiveness of different measures. While the appeal of the movement rests on an idea of experts as impartial scientists that only provide evidence to decision makers, their practice often pushes experts to provide opinions on the policy decision as a whole, thus overstepping the boundaries of the uncertainties over which they have epistemic authority. I illustrate this mechanism with the example of a scholastic policy implemented in India, which has been first tested and then actively supported by researchers.

Finally, chapter four tackles the recent COVID-19 pandemic, as a case in which severe and complex uncertainties made timely, effective policy making particularly difficult. But rather than taking the uncertainty as a reason to delegate policy making to experts, I argue that different types of uncertainty all provide reasons to make our collective decision making more inclusive, especially with respect to marginalised groups.

In conclusion: there is more to uncertainty than ignorance about how things are or will be. Uncertainty is plural, as it can concern a variety of aspects and show different properties. It can concern values and priorities, making it not entirely susceptible to being reduced with empirical methods. But while the focus of the thesis is on this plural uncertainty, there are two other recurring themes that it may be helpful to highlight here. The first is the concept of relevance: besides having its own dedicated chapter, relevance plays a crucial part in model uncertainty. What is deemed relevant has implications not just for individual, but also for collective, public, and even scientific decision making. Overlooking something important and wasting resources on confounding factors are sins of relevance, and understanding the mechanisms of relevance can only help in treating and avoiding them. The second theme is the role of reasons in decision making. Even though I do not devote a specific chapter to reasons, these are central to my accounts of both uncertainty and relevance. The importance reasons have in deliberation is undeniable – and yet, works bridging the philosophy of reasons and decision theory are still very sparse. Perhaps this thesis on uncertainty can also contribute to shed some light on the connections between relevance, reasons, and decisions.

Chapter 1

Types of uncertainty in decision making

Abstract – While mainstream decision theory only allows for variations in the severity of uncertainty, the plurality of labels with which uncertainty has been referred to in the literature and the variety of doubts that decision makers can have seem to suggest that there are different types of uncertainty. Given the importance that uncertainty has in almost any decision, understanding this plurality can be helpful to decide effectively. I propose an account of uncertainty as based on a disagreement between reasons supporting alternative mental attitudes. Under this account, dealing with uncertainty means dealing with disagreement; however, this disagreement can be radical, i.e. persistent under ideal cognitive and epistemic conditions. When this is the case, the disagreement and therefore the uncertainty cannot be resolved with an increase in evidence. I draw a typology of uncertainty reflecting the conditions that must obtain for the possibility of radical disagreement, and I trace the role that each of the types identified plays in decision making.

1.1 Introduction

Uncertainty is a pervasive feature of life. From the smallest choices to the big issues of our society, we may not be sure about what we want, what to do, and what will happen. The extent of our doubts can be paralysing, and yet in the face of all this uncertainty we still need to take action. But while decision theorists have been studying principled ways to make decisions under uncertainty, investigations of the concept of uncertainty itself are still limited.

The pervasiveness of uncertainty makes it central to many different fields, from philosophy and economics to climate science, medicine, psychology and management. This means that the concept has been addressed from different perspectives and using different labels. Discussions on uncertainty include mentions of risk, ignorance, ambiguity, unawareness, and indeterminacy, as well as distinctions between epistemic, aleatory, external, internal, fundamental, procedural, objective, subjective, ontological, normative, moral, ethical, empirical, additive, multiplicative, Keynesian, Knightian, severe, deep, great, strong, empirical, and modal uncertainty – among others.

Uncertainty is problematic because it makes it hard to choose what to do. Ultimately, we need to act effectively in our environment, and conditions of uncertainty hinder our efforts to move adequately in the world. We constantly face *practical* uncertainty (Peter 2021), or uncertainty about what to do. In order to solve it, we need principles to guide our decision making. For that, we need to understand the uncertainty we face: and if there are different types of uncertainty, as the plurality of labels in the literature suggests, then these may require different decision approaches.

Thus, the aim of this paper is to contribute to the understanding of decision making by proposing an account of the different types of uncertainty that can play a role in it. I start by presenting the way in which uncertainty is treated in decision theory (section 1.2). While covering somehow quantitatively different types of uncertainty, the default, probabilistic approach leaves out qualitative variations. To introduce this variety, I look at all the uncertainty that an agent can face when making a decision (section 1.3), which can concern all the elements of the decision model. In section 1.4 I propose an account of uncertainty as a matter of conflicting attitudes: given this account, uncertainty stems from disagreement, and dealing with uncertainty means dealing with disagreement. But this disagreement can be radical (section 1.5), and one cannot expect to solve uncertainty due to radically conflicting attitudes through empirical and logical inquiry. In section 1.6 I propose a typology of uncertainty based on the attitudes involved and the sort of disagreement in which these can incur. Finally, I go back to the variety of doubts explored in section 1.3, and suggest what type of uncertainty is likely to be at play in each case. Understanding uncertainty in terms of disagreement clarifies that not all the doubts that an agent may face are due to lack of information and can be solved with more evidence.

1.2 Uncertainty in decision theory

Imagine that you are leaving for work, and have to decide whether to walk there or take the bus. Taking the bus is faster – unless there is a traffic jam, in which case walking is safer. Decision theory would model your decision problem in a table like this one:

	Traffic	No traffic
Walk	*Arrive right on time*	*Arrive right on time*
Take the bus	*Arrive late*	*Arrive early*

In the table, the left column lists the alternative options you face, the top row the different conditions you may find, and in the other cells the outcomes resulting from performing some option under some condition. In mainstream decision theory (specifically, the expected utility tradition following more or less loosely von Neumann and Morgenstern 1944 and Savage 1954), you should then assign numerical values to the outcomes depending on their desirability and sum them for each act, weighed by the probability of the condition under which they each obtain. Then, you should choose the act that maximises that sum.

This is how standard, textbook decision theory models decisions. In this example, all the uncertainty faced by the agent at the moment of leaving for work has been repreented with probability values assigned to possible traffic conditions. The only thing that the agent has doubts about is whether there will be traffic or not; the rest is taken as given, as known. In general, probabilities are supposed to be able to represent entirely the agent's uncertainty.

The intuition that probability is a complete representation of uncertainty has limits that were famously pointed out by Knight (1921). He distinguished between measurable and unmeasurable uncertainty, where only the latter is properly uncertainty, while the former is labelled as *risk*. While in both cases the agent does not know whether some event will obtain, under risk it is nonetheless possible to make probability judgements about it in terms of statistical frequency or objective chance (Lawson 1988).

The distinction between uncertainty and risk based on the accessibility of objective probability has become standard, especially within economics. Luce and Raiffa (1957) established the usage by distinguishing between certainty, when the agent knows that the actions lead to specific outcomes; risk, when they know that the actions lead to one of a specific set of outcomes, each with given probability; and uncertainty, when they know that the actions lead to one of a specific set of outcomes, but the probabilities for each outcome are unknown or meaningless.

Savage (1954) elaborates a theory of decision under uncertainty in which, whenever the agent's patterns of choices respect some axioms, then the agent can be represented as if they assigned probabilities to states, even when no objective probabilities are given. This theory provides the grounds for a subjectivist interpretation of probabilities as measuring degrees of beliefs: these can be informed by aleatory and statistical information, but are not identical to it.[1]

Savage's result suggests that all uncertainty can be measurable: within a subjectivist interpretation, all uncertainty can be captured by a single probability function, as long as some axioms are satisfied. Thanks to the success of Savage's theory, this is still the default view in many fields, and as such it has been highly discussed. Most objections have focused on the descriptive and normative adequacy of Savage's axioms, as is the case with debates around the famous Allais' paradox. But a different question, and the one with which we are concerned here, regards not its adequacy with respect to the cases it covers, but the *scope* of application of the theory. As we try to understand the concept of uncertainty, we are interested in whether the default view can really cover all the uncertainty that decision makers can face.

A possible source of limits to the scope of the default view comes from the *severity* of the uncertainty: it may be that the information available is such that it does not warrant a representation of uncertainty with a single probability function. For instance, Ellsberg (1961) shows that there is a class of situations in which people do not act nor wish to act in accordance with Savage's axioms, so that their beliefs cannot be represented with any meaningful probability function: differently from other cases, in these situations the violation arises because of the severity of the uncertainty – the information available is not enough to support one distribution of probabilities over others. In yet more serious circumstances, no information whatsoever about the likelihood of some event may be available, or there may be events the possibility of which agents may entirely ignore.

Critical considerations like Ellsberg's have led to a refinement of Luce and Raiffa 1957's original categories. The epistemic circumstances of the decision can then be described using the following terminology, which

[1] As many other things, Savage's theory was already anticipated by Frank P. Ramsey (1926). However, it was mostly Savage's book that influenced following works and helped shape the default view. See Bradley (2004) for discussion.

provides a classification of types of uncertainty according to their severity:

(i) Risk

(ii) Uncertainty

(iii) Ambiguity

(iv) Ignorance

(v) Unawareness

Risk identifies the contexts in which aleatory and statistical information is sufficient to assign objective probabilities to the relevant events. *Uncertainty* refers to all those contexts where the information is enough for the agent to behave as if they assigned (subjective) probabilities. *Ambiguity* refers to contexts that permit only some confidence in one distribution among others, as in the situations described by Ellsberg. Under *ignorance* the agent has no information whatsoever regarding the likelihood of different events (Bradley 2017; Einhorn and Hogarth 1986), and under *unawareness* they are not even aware of their possibility (Schipper 2014; Steele and Stefánsson 2021).

As this overview shows, distinctions between types of uncertainty based on their severity date back to Knight and are now widely recognised, along with the potential limits of the default view to cover them all. However, types of uncertainty may vary not only in terms of their severity, but also of their *quality* (Bradley and Drechsler 2014): perhaps there are irreducible, different kinds of uncertainty, which cannot be covered by the default view. For instance, Hansson and Hirsch Hadorn (2016) see the default view as a "reductive" approach which disregards most types of uncertainty to make the decision treatable under a specific type of formal analysis.

In this paper, I will explore this second line of criticism, by proposing types of uncertainty that do not depend on the quantity of information regarding the actual state of the world that is available to the agent.

After all, knowledge of the state of the world in itself does not imply any choice. Even if the agent had perfect information, they may still have other doubts regarding what they should do. In the next section, I will try to illustrate what doubts the agent may face when making a decision.

1.3 The locus of uncertainty

Uncertainty over the actual state of the world is not the only uncertainty that an agent may face. Let us go back to our example on the bus/walk decision. Of course, the agent can be uncertain about whether there will be traffic or not; but they may also be uncertain about whether they should include weather considerations in their deliberation, which alternatives they have, how to evaluate possible consequences, or how reliable their probabilities are. Looking at the way in which decisions are usually represented, we can identify five elements in a decision problem:

(i) A set of available actions,

(ii) A set of possible states of the world,

(iii) A set of possible consequences of each action under different states,

(iv) A probability function,

(v) A utility function.

In order to understand the role that uncertainty plays in decision making and how to best respond to it, it is important to understand how it can enter into decisions. As long as a decision is structured in this way, then each of its elements can be the object of uncertainty – the three sets of the model, the probability function, and the utility function. Let us have a look at how uncertainty can involve elements of the decision other than the actual state of the world.

Model. Your decision about how to get to work could have been modelled in uncountably many other ways. There may be unforeseen circumstances that you did not take into account, and there may be events that you did not think were relevant for your decision: being uncertain about what the possible states are is what Bradley and Drechsler (2014) refer to as *state space* uncertainty. Our example includes the possibility of traffic but it does not include the possibility of rain, for instance, or of meeting friends on the way. One could have included those considerations as well, and obtained a more detailed model for a more thorough decision. But in fact, there are an indefinite number of things that it has not included, the majority of which are entirely irrelevant for your decision, or just not worthy of your reflection. There are many alternative models for the same decision, and these may even lead to different recommendations: the possibility of traffic may induce you to walk, while the possibility of rain may push towards taking the bus. Models that included one and not the other may yield different conclusions.

Moreover, the agent may have doubts about the identification of the alternatives at their disposal. They may be unsure about whether some option is actually available, in terms of faisibility as well as permissibility. Perhaps using a bike would be the best option, if they managed to find one, or maybe using roller-blades, if they still knew how to skate. Or perhaps they may consider whether stealing a bike to get in time at a crucial meeting could be permissible.

This means that how the decision is modelled is a normatively important matter, and that there may be better and worse models. But if this is so, how do you know whether the model you are using is in fact the one you should use? More specifically, how do you know what to include in your decision model? This is uncertainty regarding the selection of elements for the three sets of the decision model. Which possible conditions should the agent consider? Which possible outcomes should they evaluate? Which actions should they take as their options?

Utility. The agent may be uncertain about how to evaluate some possible outcome. This *evaluative* uncertainty can hit on two different dimensions: the single values and the overall function. Let me try to make this distinction clearer.

A utility function is a set of values assigned to possible outcomes. These values may be perfectly precise, but they can also be entirely unknown or only known to fall within a range. An agent with imprecise or unknown utilities will be more uncertain about their decision than one with precise values. This uncertainty may be due to doubts about whether some outcome has a certain property, and therefore about its proper evaluation. Or it can be due to the complexity of outcomes composed of a variety of aspects: the agent may like some of them and dislike some others, to the extent that it is hard to form an overall, all-things-considered evaluation. When too many things are at stake, uncertainty about the net value of the alternative is a possibility.

However, the agent may also be uncertain about two sets of equally defined values. For instance, a woman considering pregnancy may not know whether she should evaluate the prospect of motherhood according to her current utility or according to the one she would have if she had the child: this is an example of a transformative decision, i.e. a choice that could change what you value (Pettigrew 2019). Assuming that she knows both her current utilities and those she would have as a mother, then her uncertainty is different from one that only concerns the precision of values. This sort of uncertainty can also arise because the agent is uncertain about the relevant moral or aesthetic considerations: they may not know the answer to an ethical dilemma because they are unsure whether they should be utilitarian or not about it (and not because they do not know which values an utilitarian would assign), or they may be uncertain about the value of a painting because they do not know whether to evaluate it in terms of the pleasure they obtain from it or of its originality (and not because they are ignorant in art history and do not know how to evaluate the painting's originality). Finally, uncertainty over the utility function can also arise due to the presence of a plurality of stakeholders. Different people may evaluate the same outcome differently, depending on the impact it has on them: any single evaluation is bound to favour or disfavour someone. Again, this is uncertainty about whose utility should be used, and not about which values someone's utility assigns.

Probability. The agent may not know whether their probability assignment is correct. This is uncertainty about one's uncertainty: the probability assignment is supposed to reflect the agent's uncertainty over the possible states of the world, but the agent may be uncertain about the reliability of the information they have (see Hansson (1996)'s *uncertainty of reliance*), or be aware that their probability judgements are based on very poor evidential grounds.

Thus, uncertainty over probability is second-order uncertainty over states. This sort of uncertainty can be represented with second-order probabilities, or with weights over different probability distributions (e.g. Chateauneuf and Faro 2009; Gärdenfors and Sahlin 1983; Klibanoff, Marinacci, and Sujoy Mukerji 2005). A possible measure of this second-order uncertainty was already proposed by Keynes (1921) with his notion of *weight* as a measure of evidence (Dequech 2000).

Distinctions of order can go beyond the second and into a potentially indefinite progression (Dow 2012). Moreover, they are not necessarily limited to uncertainty over states. For instance, the agent may not know what to do – but they may also not know what to do about this practical uncertainty. When you face a decision, you not only face the uncertainty regarding that decision: you may also face the uncertainty regarding how to go about resolving that uncertainty. You may wonder, for instance, whether the decision is strategic or not, and therefore whether it requires game-theoretical solutions or not. Or you may face a decision concerning a plurality of stakeholders, and wonder about the right social choice procedure to address it.

This overview does not necessarily present a complete taxonomy of the uncertainties that agents may face in decision making: some of these may overlap, and perhaps others are possible. What is important, however, is that while mainstream debates tend to focus on uncertainty over possible states, *all* the elements of a decision can be the object of some uncertainty, and there is no reason to assume that this uncertainty should be treated in the same way in all cases.

As a final note, it is worth mentioning that uncertainty with respect to decisions can be uncertainty on the part of the agent or uncertainty on the part of the modeller, in case the two differ – as is usually the case with decision-theoretical models used in economics. Then, the modeller has to take into consideration all the uncertainty that the modelled agent may face, while facing uncertainties of their own, regarding for instance what they do not know about the agent or regarding the adequate theoretical tools to use. However, we are here concerned with uncertainty exclusively from the perspective of the agent, given that we are looking at uncertainty as an obstacle for effective decision making, so we will leave this distinction aside.

1.4 Uncertainty

Unitary conceptions of uncertainty tend to trace it back to a limitation in the agent's knowledge due to some lack of information (Dow 2012; Winkler 1996), and often see typologies to be at most pragmatically useful when dealing with complex issues. However, our exploration of the different *loci* of uncertainty in decisions supports the suggestion that there is a plurality of types of uncertainty, going beyond the ignorance of some relevant events that is modelled in the probabilistic language we have described in section 1.2.

When thinking about uncertainty, an obvious starting point is to take it to be the opposite of certainty. Reed (2021) distinguishes between two different senses of certainty. *Psychological* certainty is the feeling of the person that is perfectly convinced of the validity of their opinions; uncertainty would then be the feeling of having some degree of doubt.

While interesting per se and relevant to understand human decision processes, this is not the sense of uncertainty that we try to capture here. Note that, as long as one's concept of knowledge implies that what is known is true, psychological certainty does not entail knowledge, as one can feel certain about something wrong (the case that Floridi (2015) calls "insipience"). On the other hand, *epistemic* certainty is the highest epistemic status that beliefs can enjoy. In turn, this can be interpreted as a property superior to knowledge or as the highest form of knowledge: either way, epistemic certainty entails knowledge, but the converse is not true. While according to Reed (2021) we may not currently have a satisfactory account of epistemic certainty, this general picture suggests that uncertainty cannot simply be taken to be lack of knowledge, as knowledge itself can fall short of certainty.

However, Reed's distinction between psychological and epistemic uncertainty is not exhaustive, because the concept of uncertainty as an epistemic property does not include the possibility that uncertainty may concern attitudes beyond beliefs. Attitudes are standardly distinguished between *cognitive* and *non-cognitive*, where the first are those that purport to represent reality (e.g. belief, but also suspect, doubt...), while the second are those that do not (e.g. desire, hope, aversion...). Given that it seems possible to have some degree of doubt on all sorts of judgements, including those expressing non-cognitive attitudes, an account of uncertainty that managed to cover doubts on both types of attitudes would be preferable.

The question of uncertainty arises because we need to act effectively in the environment. Thus, all the uncertainty faced by the agent is problematic insofar as it hinders their ability to deal with *practical* uncertainty, or uncertainty over what they should do. If the agent is uncertain about what to do, it means that (i) there are more than one mutually exclusive alternatives, and (ii) the agent has (motivating) pro tanto reasons for all of them. If either of these conditions failed, the agent would not be uncertain about what to do: the uncertainty arises from *conflicting reasons*, that is, reasons supporting mutually exclusive alternatives [2].

Makins (2021) proposes to interpret uncertainty over non-cognitive attitudes by appealing to the psychological notion of *ambivalence*, which is the predicament of an agent having pro tanto competing reasons for alternative options. The agent has pulls in opposing directions – a situation which is different from indifference, where they have no significant pulls either way.[3] Just as uncertainty, ambivalence is a gradable notion, the degree of which is determined by the balance and weight of the competing reasons: the closer the balance, and the stronger the weight, the higher the ambivalence. Thus, for Makins ambivalence is a form of doubt that is not due to lack of information like uncertainty, but to the presence of competing reasons.

However, the opposition between uncertainty for beliefs and ambivalence for other attitudes may not be as strict. Information, evidence, theoretical considerations – these can all be reasons to believe something. The agent is uncertain about some proposition because they have reasons to believe it, but also reasons not to. When uncertain about tomorrow's weather, I have reasons to believe that it may rain – that we are in a wet season and that it rained today, for instance – and also to believe that it may not – this may be what the forecast says. If all my information points in the direction of rain, then my uncertainty will be significantly lower. I may still consider the information to be inconclusive, which on the other hand may in itself be a reason not to believe that it will rain. In other circumstances, scarcity of information itself may be a reason not to believe something, and considering something to be possible a reason to believe it may happen.

We can then generalise the intuition that uncertainty is a matter of conflicting reasons and propose the following definition:

[2] I take the absence of reasons in favour of some option to be a reason for its alternative(s); while in some contexts this could be problematic, for our purposes the simplification should be harmless and will permit us to avoid some cumbersome digressions.

[3] Makins notes that the decision-theoretical notion of indifference may cover both, as long as they result in no strict preference.

> **Uncertainty**. An agent is uncertain about some attitude $V(a)$ if and only if (i) there are some mutually exclusive alternatives to $V(a)$, and (ii) the agent has pro tanto reasons to hold each of the alternatives.

Some clarifications are due in order to understand this definition. By (cognitive or non-cognitive) *attitude* I refer to an intentional mental state, i.e. a mental state about some content. For cognitive attitudes, I assume that this content is propositional, i.e. that what is believed are propositions: that all attitudes have propositional content is not uncontroversial (Grzankowski 2015), but it is a simplification that I do not extend to non-cognitive attitudes. Finally, notice that this definition does not capture psychological uncertainty. It is possible for someone to feel doubtful even in the absence of competing reasons, just as it may be possible to feel certain even when there is some conflict. Instead, the definition aims to capture the notion of uncertainty with respect to both cognitive attitudes (epistemic uncertainty) and non-cognitive ones (Makins' ambivalence).

In the definition of attitude I include choice, which I take to be a mental state about some option; actions can be interpreted propositionally (as in Jeffrey 1965 and J. M. Joyce 1999), which means that the content of choice can be propositional, if one needs it to be. Then, practical uncertainty means that the agent is uncertain about what to choose because there are mutually exclusive alternative choices and they have reasons for all of them.

Finally, a few words about reasons. As mentioned above, the reasons involved in this account are motivating, rather than normative (Alvarez 2017; Dancy 2000b) – namely, reasons in the eyes of the agent, rather than whatever the agent should consider a reason. These can conflict in different ways: because they support alternative actions, for instance, or because they support incompatible claims (Sher 2019). Beside this, the nature of reasons has long been investigated in philosophy. They have been taken to be facts (Raz 1999), evidence (Broome 2013), and propositions (Dietrich and List 2013; Sher 2019), among other things, and their connection with internal motivations is highly debated (Korsgaard 1996; Williams 1979). In this text, I will follow Scanlon (1998) and limit myself to saying that reasons for something are considerations counting in favour of that something (p. 17). Nothing more substantial about the nature of reasons should be required by this account of uncertainty.

1.5 Disagreement

If uncertainty is a matter of conflicting reasons, then resolving the uncertainty means resolving a disagreement; and more in general, dealing with uncertainty means dealing with disagreement – albeit not between people, but between reasons. If this is so, then understanding the nature of the underlying disagreement may be necessary to properly approach decision making under uncertainty.

In general, if there is a disagreement, we may think that there is a mistake somewhere: one of the sides may be misled by some biases or cognitive shortcomings. Once these are removed, the disagreement may persist due to epistemic limitations, like lack of some information or ignorance of relevant facts. If this is the case, then reducing these epistemic limitations should eventually reduce the disagreement as well. I will call this sort of disagreement *amenable*. But in some circumstances the disagreement may survive even under ideal cognitive and epistemic circumstances. This sort of disagreement can be called *radical* (Tersman 2006), and does not tend do go away with increases in evidence or other epistemic progresses.

The nature of this distinction is such that the two types of disagreements are not approachable in the same way. While we can expect evidence and epistemic investigations in general to be able to resolve, at least in principle, amenable disagreement, the same cannot be said for radical disagreement. It is possible that some disagreement has both amenable and radical components, so that better epistemic conditions may lead to a reduction of the overall disagreement by dispelling the conflict on some aspects of the matter at hand; but the radical components will persist and will not have been reduced.

If there are different types of disagreement, and uncertainty is a matter of disagreeing reasons, then we may expect to find corresponding types of uncertainty, as long as we can expect the reasons underlying uncertainty to be in both amenable and radical disagreements. Thus, uncertainty arising from radical disagreement will not be resolvable with improvements in the epistemic or cognitive conditions of the agent, while the one coming from amenable disagreement will be. We can expect to resolve uncertainty with increasing evidence or by removing biases only if the alternative reasons are in an amenable conflict, otherwise we should expect it to persist even under ideal conditions.

What is crucial now is to understand in which cases we can have radical disagreement between reasons, if ever. As we have seen, there can be uncertainty over both cognitive and non-cognitive attitudes. A central difference between the two is the "direction of fit", which is mind-to-world for those attitudes whose content should conform to the world and world-to-mind for those whose aim is for the world to conform to their content (Björnsson and McPherson 2014). This means that while cognitive attitudes are evaluated in terms of the accuracy of their fit with reality, non-cognitive attitudes are not. For this reason, we say that a belief is true or false, but not a desire. Note that while we tend to assign truth values to cognitive attitudes, truth is more precisely a property of the proposition that the mental state contains.

Given that cognitive attitudes are evaluated on the accuracy of their content, we can expect evidence and other epistemic considerations to be relevant to their assessment. They can provide reasons to believe something, eliminate reasons in favour of some option, and ultimately settle the question of what the content of a cognitive attitude should be. On the other hand, non-cognitive attitudes are not evaluated in terms of accuracy, so the alternatives are not favoured in virtue of their correspondence with reality. If this is so, then radical disagreement is always possible with non-cognitive attitudes: some reasons may support different alternatives on grounds other than their accuracy, meaning that they would not change under better epistemic conditions.

But radical disagreement may not be a prerogative of non-cognitive attitudes, and be possible even with cognitive ones. For instance, it could arise when the proposition over which there is disagreement – i.e. the content of the attitude – is neither true nor false, perhaps because there is no corresponding fact of the matter. In that case, more evidence will not settle the disagreement on whether that proposition is true or false. However, someone who supports the classical principle of bivalence will deny that any such proposition exists, because all propositions are either true or false. And many-valued logics rejecting the principle do not provide much consolation, because rejecting bivalence does not imply that some proposition lacks a truth value: rather, it means that there are values beyond true or false. If this is so, then disagreement over the proposition could still be resolved with the assignment of a specific third value.

We do not need here to take a stance on the principle of bivalence, or even on whether all propositions have a truth value. Even if they did, this would not imply that all propositions are equally epistemically accessible. Fitch's paradox, for instance, is a challenge to the claim that all propositions are knowable (Brogaard and Salerno 2019). Moreover, moral sceptics may think that moral judgements can never be known or justified (Bambrough 2020), and the value of transformative experiences may be inaccessible *ex ante* (Paul and Quiggin 2018). If a proposition is intrinsically inaccessible, then no amount of evidence could ever settle the matter over its truth: disagreement over epistemically inaccessible propositions may be radical.

Again, someone may be sceptical with respect to the existence of inaccessible propositions (Edgington 1985; Van Benthem 2004; Van Ditmarsch, Van der Hoek, and Iliev 2012). And again, we do not need to settle the issue here. After all, we are trying to understand the notion of uncertainty *for decisions* – which means that we may considerably restrict our scope. It may be possible that there are no in principle inaccessible propositions, but that some propositions are hard enough that they cannot be known given the horizons of the decision, and can therefore be considered inaccessible for the purposes of the decision at hand. A moral realist could think that there are moral facts, that propositions expressing them are either true or false, and that these facts are accessible to our knowledge – and still believe that we are in no way close to discovering such facts, so that uncertainties over moral facts within a decision can safely be treated as if these were inaccessible, given that no epistemic investigation is likely to settle disagreements over the issue within the time horizon of the decision.

I am not contending that this is the case, and that even if there are moral facts then we cannot know them in time for our decision. I am not making claims with respect to what is metaphysically or epistemically the case; rather, I am tracing the route of how it could be possible for radical disagreement to exist. Disagreement can always be radical when concerning alternative non-cognitive attitudes; but it could arise even with cognitive ones. Disagreement could be radical when the proposition at stake does not have a truth value (if any such proposition exists), given that then no amount of evidence could settle the issue of its truth. It could also be radical with propositions that do have a specific truth value, but that are not epistemically accessible (if any such proposition exists), given that then no information could make them knowable to us. Furthermore, for the purposes of decision making disagreement could be radical with propositions that have a truth value and that are in principle accessible, but not within the horizons relevant to the decision at hand.

The fact of the matter that remains is that there are cases in which disagreement does not tend to go away with the removal of cognitive or epistemic obstacles, even with cognitive attitudes. Whatever the metaphysics or epistemology behind it, the phenomenon persists (and may even be rational under some circumstances, see Nielsen and Stewart 2020), and there is no reason why we should not expect disagreements between reasons to show the same behaviour. Moreover, admitting the possibility of radical disagreement does not lead to scepticism. The phenomenon could well be very limited in scope, far from requiring that we cannot know *anything*. Finally, it is important to note that none of these conditions *implies* radical disagreement, so that whenever they hold then there must be some radical disagreement. What they do is to allow for its *possibility*, but they are entirely consistent with lack of disagreement or with a purely amenable one.

Reasons disagree by supporting mutually exclusive alternatives. The uncertainty is resolved when there are no more conflicting reasons, that is when all the reasons the agent has support the same alternative (or their balance is tilted enough in favour of one of the alternatives for the agent's purposes). Note that if all the reasons are in favour of one alternative, but they are somehow insufficient or inconclusive, the insufficiency itself can be a reason against that alternative, so that the uncertainty could still persist. In case the disagreement is amenable, we can expect that the removal of some cognitive or epistemic limitation will move the balance of reasons towards one alternative. In case of radical disagreement, there is no reason to expect a convergence, and therefore a resolution of the uncertainty.

1.6 Types of uncertainty

In this section, I present a typology of uncertainty based on our discussion of disagreement. Thus, I will categorise uncertainties depending on whether they concern cognitive or non-cognitive attitudes, and on whether the contended matter is one over which there may be radical disagreement.

Empirical and logical uncertainty. Cognitive attitudes purport to represent reality and their content is apt for truth value, which they inherit. But it is important to remind that aptitude to truth values may not necessarily imply that the truth value is determined or accessible. We can therefore distinguish between determinate and indeterminate propositions, where the second are those whose truth value is not available to a perfectly informed agent [4]. The value of determinate propositions can usually be settled with evidence: there is some amount of information that could, at least in principle, solve the agent's uncertainty over their value. Reasons provided by evidence tend to support the true value of the proposition. This is *empirical* uncertainty. Alternatively, it can be settled with logic: tautologies and contradictions have determinate truth values independently of empirical evidence. Uncertainty over them would be *logical* uncertainty. Given the sensitivity to evidence and cognitive conditions, disagreements over these propositions will be amenable. However, the same cannot be said for propositions with indeterminate value (Pravato 2020).

Vague uncertainty. Indeterminate propositions can be of two types as well. The first type are propositions with *vague* concepts (Sorensen 2018). These are concepts that have borderline cases, as is the case for instance with the predicates "old" or "child": there is no specific age benchmark after which one is old or stops being a child. For this reason, there are uses of the concept "old" that make some agent legitimately uncertain about their value. That someone is not young may be a reason to take them to be old, while lack of grey hair may be a reason to the contrary, thus generating the uncertainty. Note that no amount of conceptual analysis or empirical research can settle the question. We could stipulate that "old" starts right after one passes the middle of the average life expectancy, but we would be creating a new concept: considering someone young one day and old the day immediately after does not correspond to our shared understanding of the concept "old". Uncertainty concerning the status of borderline cases of vague concepts I call *vague* uncertainty [5]. It is possible that disagreement concerning border-

[4]This can be because the truth value is somehow missing or it is simply inaccessible, depending on one's assumptions.

[5]Eriksson and Olinder (2016) call this *conceptual* uncertainty, but I find the term confusing, as it is not uncertainty over the meaning of a concept but over whether

line cases persists even under ideal epistemic and cognitive conditions, and so will the uncertainty.

Ontic uncertainty. The second type of indeterminate propositions concerns those propositions the truth value of which depends on some non-deterministic aspects of reality. I will call the corresponding uncertainty *ontic*. In the literature, this type of uncertainty has long been recognised in opposition to uncertainty as a property of the agent (see e.g. Davidson 1996; Dequech 2004; Fishburn 1994; Kahneman and Tversky 1982; Kozyreva and Hertwig 2021; Perlman and McCann 1996).

Even though distinctions on these lines are quite widespread in the literature on uncertainty, appearing under labels like epistemic/aleatory or internal/ external, the traditional framing in terms of a reality/agent opposition is confusing: as we have defined it, uncertainty is always a property of the agent, not of the world. The relevant distinction concerns the *source* of the agent's uncertainty. Ontic uncertainty does not mean that the world is uncertain, but that the agent is uncertain due to the world being indeterminate – that is, of course, assuming some indeterminacy in the world, which could be due for instance to the actions of sentient beings or to the non-deterministic nature of the world that seems to be suggested by quantum physics.

The relevant aspect of ontic uncertainty is that it is about some state of the world that does not exist yet, which means that there is no amount of information that can solve the uncertainty before the relevant state obtains. Notice that the fact that the event has not happened yet does not mean that we know nothing about its possibility. We may have sufficient information to assign precise probabilities to its occurrence, and if we do not, we may look for it and thus reduce our uncertainty. We may have accurate information regarding the chance that a certain quantum event happens, and if we do not, we may want to ask experts for their opinion or study the phenomenon further. The uncertainty over the probability of the event may be empirical, but the uncertainty

some objects are instances of that concept. With vague concepts, this can happen even with perfect mastery of the concept at issue.

over the event itself is ontic: we have (possibly unbalanced) reasons to believe both that it will happen and that it will not, which means that we have conflicting reasons for two alternative contents of our belief and, consequently, we are uncertain about it. This uncertainty cannot be resolved with evidence.

Of course, that there is genuine indeterminacy in the world is an assumption that many may decide not to share. The debate on whether free will exists and, if so, if it is compatible with determinism is well alive in philosophy (see O'Connor and Franklin (2021) for an overview and List (2019) for a recent contribution), and so is the one on the meaning of quantum mechanics in physics: while some take it to show that there is some genuine chance in the physical world (e.g. in the modal tradition: see Lombardi and Castagnino (2008)), others take it to be compatible with a deterministic world (e.g. in the Everettian tradition: see Albert and Loewer (1988) and Saunders (2010)). The possibility of adopting a fully deterministic view of reality is open. In that case, ontic uncertainty would be reduced to empirical uncertainty. In general, the scope of issues over which one can have ontic uncertainty depends on the extent of one's determinism.

Psychological uncertainty. Let us now move to uncertainty about non-cognitive attitudes. Tastes and emotions can be considered to be non-cognitive: basic emotive assertions like "ouch" or "yuck" are not taken to express beliefs.

One could wonder whether uncertainty over tastes or emotions could be possible at all. It seems that one could have nuanced emotions, or feel more or less strongly about something – but there is no sense in which it is possible to be uncertain about one's own emotions or tastes: any doubt should be solved with a little introspection. However, given our account of uncertainty, we can ask whether it is possible to disagree over tastes. Let us look at the following exchange:

A: Coffee is tasty.

B: No, it's not. It's too bitter.

There is a sense in which the exchange is an instance of disagreement, as the two agents disagree over whether coffee is tasty or not. However, it seems that neither is at fault, in the sense of making a mistake with their position: we would then have a case a faultless disagreement (Huvenes 2014), a concept that some view as intrinsically contradictory (e.g. Iacona 2008). If no disagreement is possible, then no uncertainty is possible. Nonetheless, while we may not have interpersonal disagreement over tastes, the situation may be different for *intra*personal disagreement: two people may not disagree, but the reasons for someone's attitudes maybe can. I can enjoy the smell of coffee, for instance, while disliking its bitterness; I can like the taste of some food, while finding its texture offputting. These are conflicting reasons: smell is a reason to like coffee, while bitterness is a reason not to. In cases like these, my attitudes are *ambivalent* in Makins (2021)'s sense, and thus I can be uncertain about my taste for coffee because I have conflicting attitudes towards it. Assuming that I have a thorough experience of coffee and no taste dysfunctions, then we can assume that this disagreement and the consequent uncertainty will persist even under ideal conditions.

Moral uncertainty. Another set of attitudes are the ones expressed by moral statements. Disagreements and uncertainty about moral issues seem to abound. People often disagree about what is good or right, and this disagreement seems to be importantly radical. The persistence itself of philosophical debates that do not seem to be moving towards some consensus is evidence of how disagreement over moral questions is resistant to empirical inquiry and expertise. Consequently, moral uncertainty cannot usually be entirely resolved with evidence. Stylised examples are provided by philosophical thought experiments like the famous Trolley Problem (Foot 1967). In these scenarios, all the relevant information is provided to the reader, but the uncertainty (and the debate) over the moral thing to do is not resolved.

Whether moral judgements are cognitive or not is a very open debate (see van Roojen 2018). While moral cognitivists claim that moral considerations are apt for truth values, non-cognitivists deny that, and claim that moral statements express attitudes other than belief. Non-cogitivism is a version of anti-realism about ethics, as it implies that there are no moral properties or facts to have beliefs about. But cognitivists can be anti-realists too: error theorists claim that while moral claims do express beliefs, they are all false, because the objects they are about do not exist. On the other hand, moral realism contends that there are objective moral truths, and that therefore moral considerations purport to represent aspects of reality.

Moral realism is not sufficient to put moral uncertainty on the same level with empirical uncertainty (Shafer-Landau 1994). Assuming the existence of moral properties and facts, taking them to be part of our environment, does not necessarily imply that we are also able to know them in the same way we can know other aspects of reality. If we cannot, then moral uncertainty, while being cognitive, would be intrinsically different from empirical uncertainty because its objects are in principle unknowable. Note that moral skepticism of this kind is open to all main metaethical positions: if we take the mainstream view of knowledge as justified true belief, we see that moral knowldege is excluded for the non-cognitivist who takes moral claims not to express beliefs, and for the error theorist because moral claims are never true. But even the realist assumption that there are true moral beliefs does not exclude the possibility that these true belief can never be justified (R. Joyce 2021).

Without taking a stance on these complex debates, the discussion shows that the possibility of non-cognitive uncertainties depends on one's metaphysical and metaethical commitments. One can be a full cognitivist by assuming that moral considerations are cognitive and that uncertainty over non-cognitive attitude is not possible. But even in that case, that is still not sufficient to justify treating moral uncertainty in the same way as empirical uncertainty, because the epistemology of moral and empirical considerations may still be significantly different. Be it as it may, widespread disagreement and uncertainty over moral questions seem to be the norm. This means that, even if there were moral facts and these were accessible, they would probably not be so within the time horizon of a decision, given that metaethical debates are far from settled, so that for decision making purposes they can be treated as if they were inaccessible.

While I have here discussed moral judgements, much of what I have said can be applied to all normative considerations: judgements on aesthetics or rationality, for instance, may be taken to express cognitive or non-cognitive attitudes and to be epistemically accessible or not – in any case, they tend to be questions over which we expect disagreement to persist under ideal circumstances, at least for the horizons relevant to much decision making.

1.7 Uncertainties in decisions

At the end of the discussion in section 1.6 we are in the position to draw the following map of the different types of uncertainty:

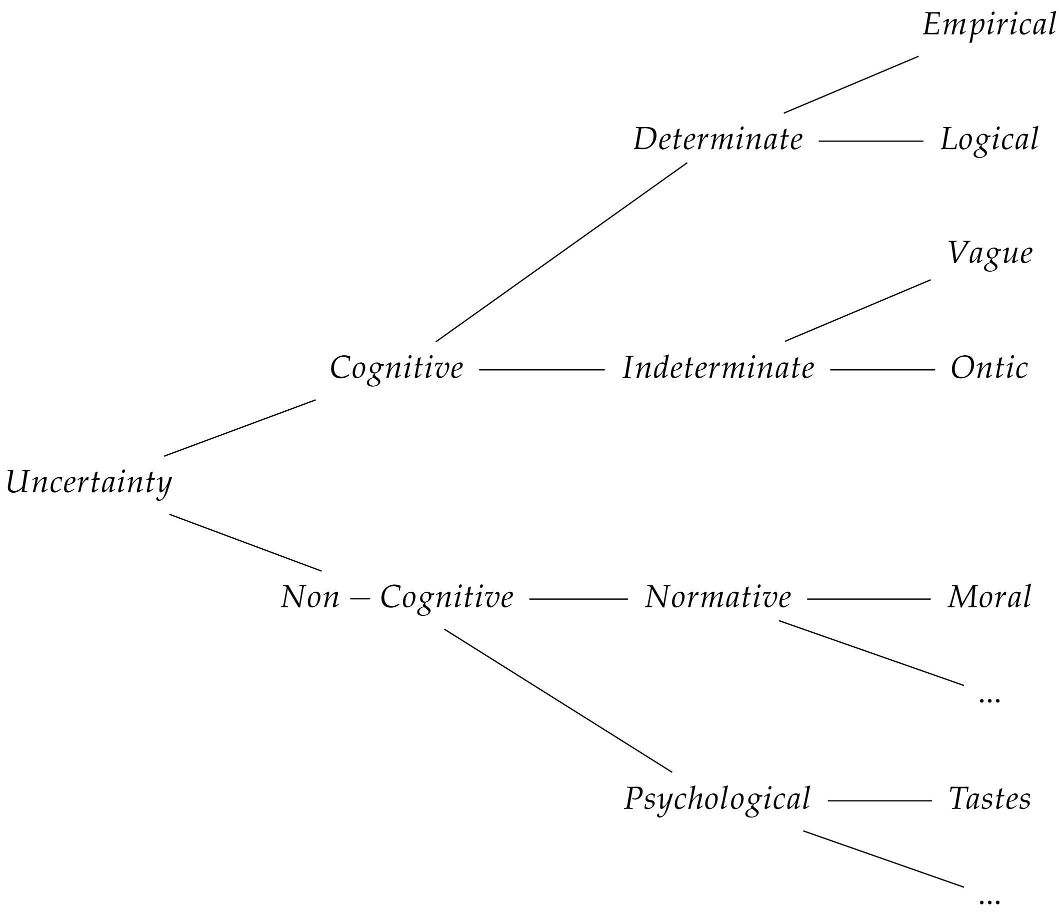

Bear in mind that the map has mostly an illustrative function. The non-cognitive category is *potential*: it includes all the cases in which the object of uncertainty may be understood to be some non-cognitive attitude, but as we have seen the debate on whether some or all of these categories are indeed non-cognitive is open. Similarly, ontic uncertainty would fall under the determinate category under a completely deterministic view of the world. Finally, the sub-categories of both normative and psychological uncertainty are open, as there may be normative considerations beyond moral ones and psychological non-cognitive attitudes beyond tastes. In the following, I will use these labels with these qualifications in mind, as shortcuts to the relative discussions in section 1.6. Among these types of uncertainty, those that can persist even under ideal epistemic conditions are the vague, the ontic, the normative, and the psychological.

We can now try to categorise the uncertainties we have identified in section 1.3.

Uncertainty over the actual state of the world – i.e. uncertainty over what is the case or will be the case – concerns the agent's beliefs, and is therefore a cognitive type of uncertainty. There are mutually exclusive alternatives, each of which is supported by some (evidential, theoretical) reasons. The uncertainty may concern empirical aspects, but also logical propositions or the status of borderline cases of vague concepts, as well as indeterminate aspects of the world: the uncertainty can therefore be of all the cognitive types analysed above. Second-order uncertainty over probabilities is mostly cognitive as well, as it regards beliefs over beliefs; however, the source of the uncertainty over these second-order beliefs may be uncertainty over the selection of experts, an issue which may include normative considerations over which there may be radical disagreement.

Evaluative uncertainty over the utility function will be normative, whenever the agent does not know which function to use in their decision making because they are uncertain about the relevant moral, aesthetic, or broadly evaluative considerations. Then, uncertainty over utility values is uncertainty about the value that the utility function employed assigns to some consequence. It can have cognitive elements – the agent may be unsure about some factual aspect of the consequences they are evaluating, for instance – but primarily it will concern psychological and/or normative aspects. The first ones will be relevant if the uncertainty comes from doubts concerning the agent's own tastes, while the second ones if it comes from doubts on what different normative considerations say about the alternative consequences.

Let us now turn to model uncertainty. As we have seen, this is uncertainty regarding the selection of elements to include in the decision model. Uncertainty over the selection of acts can involve both cognitive and non-cognitive attitudes. If it is a question of the act's *feasibility*, then it will be a primarily empirical and in general cognitive concern. If, instead, it is a question of the act's *permissibility*, then the concern will be mostly normative. The question can even have psychological traits, as the agent may consider discarding some otherwise available act on the grounds of it being too unpleasant to perform. As for uncertainty over the identification of states and consequences, it is importantly uncertainty over what matters to the agent: the consequences of the alternative options that should be considered are those that the agent wants to avoid or achieve, and the states to include in the model are those required to bring about those consequences. The agent can thus be uncertain about the model because they are uncertain about what matters to them in the decision, which would be a primarily non-cognitive uncertainty, or because they are uncertain about which conditions are required to obtain a certain outcome, which instead would be primarily cognitive.

1.8 Conclusions

Arguably, any decision happens under some degree of uncertainty: the agent can have doubts concerning a variety of aspects of the decision, and each of these can make them unsure about the best course of action. And when dealing with uncertain decisions the obvious thing to do seems to be to try to reduce the uncertainty. Understanding the nature of the uncertainty can help in doing so effectively.

The default view in decision theory takes uncertainty to be entirely representable with a single probability function. But while the debate on the limits of this view has been long and lively, it has focused primarily on cases where the uncertainty is too severe for probabilistic representations to be justified. However, uncertainty may vary not just in degrees of severity: there may be cases in which the default view fails not because the uncertainty is too severe, but because it is not of the right type. The variety of doubts that an agent can have in a decision seems to confirm the plurality of types of uncertainty. Beyond uncertainty over the actual state of the world – which is the type of uncertainty classically represented with probabilities – agents can be uncertain over all elements of the decision, from the way they modelled the problem to the utility and probability functions employed. However, this plurality does not mean that there is no unitary concept of uncertainty.

Expanding on Makins (2021), I have proposed that someone is uncertain about some (cognitive or non-cognitive) attitude whenever they have motivating, pro tanto reasons for mutually exclusive alternatives. This means that uncertainty arises from disagreement between reasons. In some cases, this disagreement can be resolved with an improvement in the cognitive or epistemic conditions of the agent: removing biases or learning something new can change the set of reasons to the point that there is no disagreement. However, in other cases the disagreement is radical, in the sense that it persists under ideal conditions. When the uncertainty arises from radically disagreeing reasons, we cannot expect evidence to resolve it.

I have argued that radical disagreement is always possible over non-cognitive attitudes, and it may even be possible over cognitive attitudes given some epistemic or metaphysical assumptions. The type and the content of the attitudes generate a typology of uncertainties, some of which will (at least in principle) be sensitive to evidence, and some of which will not. All of these types can be at play in a decision, which means that not all the uncertainty faced by an agent will necessarily be reducible with empirical evidence.

This typology is not in contradiction with other proposals in the literature. Different typologies have focused on other variable components of uncertainty: for instance, Bradley and Drechsler (2014) talk of uncertainty varying in severity and in nature, while both Hansson (1996) and Hansson and Hirsch Hadorn (2016) provide non-exhaustive lists of possible uncertainties that have been ignored by the mainstream debate. For all the uncertainties that these typologies present, one can ask what type of attitude they concern and with which content, and thus whether the disagreement at their root is radical or not. This is useful because it tells us on which of their different doubts the agent can work in order to reduce the uncertainty, and which of their uncertainties may persist even under ideal conditions.

Moreover, uncertainty is a complex notion that presents both unitary aspects and significant variability; and yet, attempts to keep both these characteristics have been surprisingly sparse, given the importance of uncertainty in our life and in our decisions. In this proposal, the variability of uncertainty has been directly connected with a unitary account of uncertainty. In turn, this account connects the discussion over types of uncertainty in decision making with fields as diverse as the philosophy of reasons, the nature of disagreement, and the debates on cognitivism. In discussing the metaphysical and epistemic assumptions of my account, I have presented some of the connections that uncertainty has with more general issues in philosophy. Uncertainty is, after all, a pervasive feature of life.

Chapter 2

A notion of relevance for rational decision making

Abstract – Decision theories have largely ignored the step of decision making in which the agent models the situation. Given that a decision can be represented with different models, and that these can lead to different recommendations, then without a principled way to assess them the agent's choice is under-determined. As models require the agent to select the aspects that matter to the decision, an account of rational decision modelling must include a notion of relevance.

I propose that the most rational model is the one taking into account all and only the considerations relevant for the decision. I define relevance for a decision as a matter of providing reasons for some option, and I identify four functional types of reasons leading to four corresponding types of relevance. I focus on what I call "constitutive relevance", which provides the content of the decision model, and propose a formal definition of this concept.

2.1 Introduction

Imagine that you are leaving for work, and have to decide whether to walk there or take the bus. Taking the bus is faster, unless there is a traffic jam, in which case walking is safer. Decision theory would model your decision problem in a table like this one:

	Traffic	No traffic
Walk	*Arrive right on time*	*Arrive right on time*
Take the bus	*Arrive late*	*Arrive early*

Table 2.1

In the table, the left column lists the alternative options you face, the top row the different conditions you may find, and in the other cells the outcomes resulting from performing some option under some condition. In mainstream decision theory[1], you should then assign numerical values to the outcomes depending on their desirability and sum them for each act, weighed by the probability of the condition under which they each obtain. Then, you should choose the act that maximises that sum.

Decision theory helps you to identify the rational decision to make, according to a certain notion of rationality. What it does not do, however, is tell you why your decision should correspond to Table 1 and not to any of these other options:

	It rains	It's sunny
Walk	*Get wet*	*Enjoy the sun*
Take the bus	*Stay dry*	*Miss out on the sun*

Table 2.2

[1]In the paper, by standard or mainstream decision theory I refer to Expected Utility Theory and its main variations. This is for ease of illustration, but most of my points should be general enough to apply to different decision theories.

	Traffic Rain	Traffic Sunny	No traffic Rain	No traffic Sunny
Walk	*On time* *Wet*	*On time* *Enjoyable*	*On time* *Wet*	*On time* *Enjoyable*
Take the bus	*Late* *Dry*	*Late* *Miss out*	*Early* *Dry*	*Early* *Miss out*

Table 2.3

	Find \$5 in your pocket	No money in your pocket
Walk	*\$5*	*\$0*
Take the bus	*\$5*	*\$0*

Table 2.4

These different ways to frame the same problem – whether to walk to work or take the bus – may lead to different recommendations. So which one should you use to ground your decision? Decision theory does not help you much there. All it can do is computing the best option for each of these frames, but it cannot tell you which one is better than which other.

This example shows a limit to decision theory as the study of rational decision making: it ignores entirely a fundamental step of the reasoning process, namely the step of decision modelling. In this paper, I will not criticise the contents of standard decision theory. Instead, I will bracket any debate about its claims and focus on something that it is simply left out. While decision theory studies the rational way to make decisions given a certain frame, I will investigate the rational way to make that frame.

Tables 1-4 provide us with a starting intuition regarding the question of which frame you should use to make the walk/bus decision. It seems that you should use Table 1 in case you only care about punctuality, Table 2 if you only care about the weather, Table 3 if you care about both, and you should not use Table 4 because it is unrelated with the matter at hand. In short, it seems that you should construct your decision problem including the aspects of the situation you deem *relevant* for your decision, and leaving out all that is irrelevant. But what does it mean to be relevant for a decision? Trying to sketch an answer to this question is the point of the present paper. This is a normative enterprise, but one which will have some descriptive implications as well.

The paper develops as follows. In the next section, I try to clarify the nature of the problem by addressing some possible doubts. In section 2.3 I set the formal framework within which my analysis develop, and in section 2.4 I provide a starting account of the notion of relevance for decisions that takes relevance to be a matter of providing reasons. In section 2.5 I identify four different types of relevance for decisions depending on the reasons they provide and the function they play, and then finally in section 2.6 I provide a more precise definition of the types of relevance that are involved in decision modelling. I conclude by suggesting some further lines of research.

2.2 Why should we care?

Some authors have vindicated the silence of decision theory on decision modelling, contending that it lays outside of the scope of the field: decision theory is concerned with the rationality of choices *given* a certain well-formed model (Steele and Stefánsson 2020). For instance, Mamou (2020) claims that decision models come with a presumption of relevance – whatever the agent or the external observer has identified as relevant *is* relevant. If different aspects of the situation were included, one would simply obtain a different decision problem: there is no "right" model of a certain situation, but one solution relative to each model. The scope of decision theory starts *after* the decision has been modelled.

However, there are a number of reasons why decision modelling may be interesting for decision theorists. First, one of the central goals of standard decision theory is to provide operational definitions of reasoning attitudes. One could describe the history of standard decision theory as the increasing introduction of subjective components in decision problems (Steele and Stefánsson 2021): von Neumann and Morgenstern (1944) introduce subjective values of outcomes, while keeping the rest objective; Savage (1954) and Jeffrey (1965) go further, by making probabilities subjective as well. However, they all leave the decision model as objective: the agent faces a decision the states and outcomes of which are given. Approaching the modelling step of the reasoning process is integral to decision theory as the next step in its subjectivisation.

Second, decision modelling is normatively salient. Different models can produce different solutions. Focusing on some aspects rather than others of a certain situation may lead to different attitudes (J. M. Joyce 1999). Your opinions towards eating veal may be different depending on whether you focus on taste or on ecological sustainability, or on both. In some cases, different grains of specification may lead to different probability assignments (Savage 1954; Shafer 1986). But the situation that the agent faces is one, and one is the choice to be made at the end of the deliberation. As long as alternative models can lead to alternative solutions, then decisions are under-determined – and of course, it is possible that one recommendation is in some sense worse than some other. Beside being normatively salient, this under-determination has descriptive implications as well: some agent's observed behaviour may be inconsistent with the predictions of a theory simply because the theorist and the agent have modelled the situation differently.

Finally, some open issues in decision theory are strongly related to decision modelling, and could therefore benefit from its study. An illustrative example is provided by the long-standing debate regarding the identification of the consequences of alternative options. As the re-description of outcomes is a standard strategy out of paradoxes like Allais' (e.g. Loomes and Sugden 1986; Pope 1995; Weirich 1986), some authors have questioned the legitimacy of this move, calling for criteria to halt the slippery slope that would lead to justification of any pattern of preferences (Broome 1991). The intuition then is that some ways to describe consequences are legitimate, and some are not. This is a question regarding the proper modelling of a decision.

If different models may lead to different recommendations, it seems that the only way to be sure that one is making the right choices is to model the decision in the most comprehensive way possible. Sure making a decision on the basis of Table 3 is safer that basing it only on 2 or 1. The best decision must be the one that takes into account every aspect of the situation. Why do we need to think about relevance?

Making a fully comprehensive model is, of course, unfeasible from a descriptive point of view: no agent can ever attain the level of representative accuracy required to be sure that *nothing*, not even the slightest detail, has been left out. But this solution is unsatisfactory even from a purely normative perspective.

The old principle that "ought implies can" means that normative claims have to take into consideration the limits of what is possible, and processing such an amount of information is not. But even if it were possible, or even if we were not interested in the feasibility of our normative claims, the most comprehensive model is not necessarily also the best model. A mirror-like representation of some situation would be useless, if the representation is to work as a model and allow surrogative reasoning (Frigg and Nguyen 2017). But what is perhaps more important, is that such a model would not be *efficient*. The additional benefit gained from the inclusion in deliberation of aspects only remotely connected with the issue at hand is not worthy of the processing effort required. While this may be undesirable *per se*, it is useful to keep in mind that decision theory is in the business of practical rationality and, more specifically, it usually aims to capture the meaning of *instrumental* rationality. Instrumental rationality dictates to take appropriate means to one's ends (Thoma 2019), and thus condemns inefficient uses of resources. Thus, instrumental rationality requires that, in order to be the best means to an instrumentally rational decision, a model should be efficient. According to this criterion, the best (most rational) decision model is the one that includes all and only the aspects that are *relevant* for the decision, so as to not overlook crucial issues while also avoiding to waste resources in pointless processing.

2.3 Setting the framework

The problem of the adequacy of a theory that only applies to already formed models to solve pre-modelled real-world decisions was already acknowledged in Savage (1954). Following him, we can identify the "grand world" (GW) model of a decision as the most fine-grained representation of the situation, including every minute aspect of the world. Any other model is a coarsening of the GW, and is called a "small world" (SW) model. Both Table 1 and 3 represent SW models, even though of different grains. The model that included all details, from the chemical composition of buses to weather in Islamabad, is the GW model. While the actual world where the agent lives is fully represented only by the GW model, decision theory only applies to SW models – which means that, for the theory to be applied, the GW model needs to be coarsened adequately.

I have argued that using the GW model is neither descriptively realistic nor normatively compelling. This means that agents both do and should construct SW models of their decisions. At the end of last section, I have suggested that the best decision model is one that includes all and only the aspects of the situation that are relevant to the decision. Following J. M. Joyce (1999), I will call this the "fully-considered" model (FC)[2]:

> **FC1** The *fully-considered* model is the model that includes all and only the aspects of the agent's situation that are relevant for the decision.

[2]To be precise, Joyce conflates the grand world and the fully-considered model. However, it is clear that the model including all and only what is relevant is not the same as the one including everything, regardless of relevance.

This is not, however, necessarily the model that agents actually do use. We therefore have two different questions of decision modelling: the normative one, regarding the criteria to assess different models of a same decision, and the descriptive one, regarding how agents do model the decision. Both involve the construction of SW models that coarsen the GW model, but the two questions are distinct. Here, I am addressing the first one.

The claim that the FC-model is the best (or at least the most instrumentally rational) model of a decision requires a good deal of clarificatory work to be made precise. What does it mean for something to be a relevant aspect of the decision?

First of all, I call *aspects* of a state of affairs all the propositions that are true of that state, i.e. all the different ways in which a state can be described[3]. The situation faced by the agent at the time of decision is a state of affairs, of which some propositions are true. The situation resulting from the performance of some action is another state of affairs, one that again presents a variety of aspects. Making a decision can be seen as choosing the option that is expected to lead to the best situation given the current situation. Some aspects of the current situation will be crucial in determining the outcome situation, and some will not. In the same way, some aspects of the possible outcome situations will be crucial to their evaluation, and some will not. When modelling the decision, the agent should include all and only those aspects of the current and the outcome situations that are crucial.

To represent this, I will take a *decision* D to be a pair $D =< A, \Omega >$. A is the set of available acts, which I take to be fixed. We can think of acts as propositions that can be made true at will (Jeffrey 1965). Taking the set of actions as given means that the agent knows her options, and has to construct a model to decide between them. This is, of course,

[3] In this paper, I adopt a propositional language to talk about the components of decision models. While this practice is shared by some decision theorists (most notably by Jeffrey (1965) and J. M. Joyce (1999)), it is by no means common ground, especially outside of philosophy. I do that for tractability, assuming that a parallel analysis can be forged under different formal frameworks.

a simplification of the actual situation. The identification of the available options is itself an important part of the decision making process, and one that should be open to scrutiny just as much as the rest of it. This simplification allows to take the investigation of relevance off the ground, but ideally it will be possible to extend the account that comes out of it to the selection of relevant options as well.

Acts are assumed to be mutually exclusive and jointly exhaustive, which means that A is a partition of Ω. Ω is the set of possible worlds described to the highest detail. Given Ω, we can define its power-set $\mathcal{P}(\Omega)$, which is the set of all propositions (intended as sets of possible worlds); $\Pi(\Omega)$, which is the set of all its possible partitions; and $\Pi(\Omega)$'s subset $\Pi_2(\Omega)$, which is the set of all its possible bipartitions. Each aspect p identifies a bipartition made of p and its complement proposition $\bar{p}$, $P = \{p, \bar{p}\}$.

I take a *decision model* to be the triple of sets of propositions $M =< A, C, S >$. A is the set of acts in D. C is the set of consequences, i.e. the set of propositions describing the situations that would result from the performance of the acts under alternative conditions. S is the set of states, i.e. the set of propositions describing these alternative conditions.

The agent assigns probabilities and utilities to states and consequences respectively: there is therefore a probability function P defined over the set S and a utility function U defined over the set C. Then, every element $a \in A$ gets an expected utility $eu(a)$, calculated by summing the utility of all its consequences weighed by the probability of the state under which they happen.

The calculations of decision theory take as input a decision model alongside a probability and a utility functions. The problem of relevance concerns the identification of the aspects of the agent's situation that should be included in M. With this framework in the background, in the next section it is time to move to the notion of relevance.

2.4 Relevance and reasons

A standard Bayesian criterion frames relevance in terms of a relation between evidence and hypotheses. In its most basic form, it affirms that a piece of evidence E is relevant for a certain hypothesis H in relation to some body of knowledge K if and only if $P(H|E\&K) \neq P(H|K)$ (Krämer 2017). This criterion has been widely debated and sophisticated in the formal epistemology literature, especially in relation with confirmation (Crupi 2021). And the notion of relevance is central to many other fields, from information theory (Borlund 2003; Floridi 2008) and law (Lempert 1976; Pardo 2013) to linguistics (Sperber and Wilson 1986) and logic (Mares and Meyer 2017). However, there are not many unified, general works on relevance.

Cohen (1994) provides a philosophical account of relevance with a rather wide scope. The core of his proposal is to interpret relevance in terms of providing reasons for something. More specifically, he defines relevance as follows:

> A true proposition R is (...) relevant to an askable question Q if and only if there is a proposition A such that the truth of R is or would be some reason, though not necessarily a complete or conclusive reason, for anyone's accepting or rejecting A as an answer to Q. (1994: 178).

The construction of a decision model is ultimately aimed at the choice of a certain option. Thus, it is reasonable to assume that something should be included in the model to the extent that it provides reasons for some option, so as to provide arguments for the agent's deliberation. We can then apply Cohen's intuition to decisions, and obtain the following definition of relevance for decisions:

> *Relevance for decisions*. An aspect $p \in \mathcal{P}(\Omega)$ is relevant to a decision D if and only if there is some available option $a \in A$ such that p provides or would provide a reason, though not

> necessarily a complete or conclusive one, for or against the choice of that option.

Now, applying this definition to models, we obtain that a FC-model is one that includes all and only the aspects of the agent's situation that provide reasons in favour or against some of the options available. In order to have a full grasp of this claim, we now need to understand what it means to provide reasons for or against an option.

Intuitively, reasons play an important role in deliberation, providing arguments to settle the indecision between some alternatives. Practical reasons – or reasons for taking a certain action – have been the object of long-standing debates in philosophy (see e.g. Korsgaard (1996), Schroeder (2007), and Williams (1979)). But important as it may be, the connection between reasons and decision making has gone widely unexplored. It is only recently that two different reason-based decision theories have started to draw more attention to it. Dietrich and List (2013) construct a preference relation over alternatives that can be represented by a weighing relation ranking different combinations of reasons. Sher (2019) shows that, under some conditions, choosing by weighing reasons is structurally identical to choosing as if by maximising expected utility. While the two results provide important instruments to bridge the gap between discussions on reasons and on decisions, neither provides a clear definition of what reasons for choice are. Sher (2019) only commits to the claim that reasons are true propositions: the weight of a reason on a decision is the effect that learning that it is true has on the decision. Dietrich and List (2013) intend reasons as propositions that are relevant for the agent's preference towards the alternatives of which they are true, but this intuitive characterisation based on relevance is circular for our purposes, as we are trying to define relevance in terms of reasons.

Both do, however, help in making the above definition *comparative* with their focus on weighing reasons: we can say that an aspect p is more relevant for a decision the heavier the reason it provides for some option. A precise cashing out of this suggestion depends on one's notion of weight of a reason. If, for instance, one follows Sher (2019) and takes the weight of a reason to be the effect this has on the decision, then the aspect providing the reason will be more relevant the larger the effect that its consideration has on the decision. In 2.6 I will make this notion of comparative relevance more precise.

The notion of relevance is crucial in an account of rational decision modelling. In this section, I have suggested that this notion of relevance could be understood in terms of reasons for options, contributing to the growing interest on the connections between practical reasons and decisions. However, a clear understanding of what these reasons are is still lacking. In this paper, I will not assume any specific view on the nature of reasons (see Broome (2013), Dancy (2000a), Raz (1999), and Scanlon (1998) for some such views), as long as reasons are such that propositions can in some sense provide them. Leaving the metaphysics of reasons aside, I will investigate their functional character, since my interest is in the role they play in decision making. In the next section, I present four different ways in which propositions can provide a reason for some option, depending on the function that the reason has in the decision.

2.5 Functional types of reasons

Let us go back to our starting example. Should you take the bus or should you walk to work? Here are four possible reasons for walking:

1. I could arrive late.
2. Being late is despicable.
3. There may be a traffic jam.
4. The radio said there has been a car accident on my route.

They are reasons for walking in the sense that each of these could appear in an agent's explanation of a preference for walking, completing a sentence like "I will not take the bus because x", or in someone's recommendation for walking, completing a sentence like "You should not take the bus because x" (with the appropriate pronoun changes; for reasons as parts of explanations see Broome 2013). However, each of these seems to provide a different type of reason: 1 provides a reason for walking by stating a possible consequence of taking the bus; 2 is a reason to evaluate that consequence negatively; 3 describes the external condition for the obtaining of the consequence; and 4 provides evidence regarding the probability of that condition.

Of course, they are not alternative reasons – one could provide a more complete explanation or recommendation by listing all of them. But each of them would be playing a different *function* in that explanation or recommendation. In order to provide a reason for an option, a proposition can state the possible consequences of the option, or the conditions under which those consequences happen. It can provide reasons to expect those conditions, or to evaluate the consequences in a certain way. The first two I call *constitutive* reasons; the last two I call *epistemic* and *evaluative* reasons respectively.

Epistemic reasons like 4 provide information about the actual state of the world and what to expect from the performance of some action: 4 provides evidence about how things are. Their function is to help define the probability function P, and they are *epistemically relevant* to the decision. Evaluative reasons like 2 are moral, aesthetic, or other considerations that determine how the agent evaluates the consequences of their actions: 2 provides a moral reason for why one should avoid the consequence of being late. Their function is to help define the utility function U, and they are *evaluatively relevant* to the decision.

Constitutive reasons constitute the content of decision models: they provide the propositions that enter into formation of the sets S of states and the set C of consequences. They are the aspects that are taken into consideration during the stage of decision modelling where the agent analyses the situation into options, possible outcomes, conditions for those outcomes. Constitutive reasons are *constitutively relevant* to the decision, and they can be so in two different ways: they can be relevant *for the set of consequences* (CRC) like 1 or *for the set of states* (CRS) like 3. The four different types of relevance, with the function they play in decision making, are listed in Table 5. We say that some aspect is evaluatively relevant for a decision, for instance, if it provides evaluative considerations that help define the utility of some option. We say that some aspect is constitutively relevant as a consequence if it is a possible outcome that has to be included in C.

Name	**Content**	**Function**
Constitutive relevance (cons.): CRC	Possible outcomes of the actions	Provide the elements of C
Evaluative relevance	Moral, aesthetic or other evaluative considerations	Help define U
Constitutive relevance (states): CRS	Conditions for the obtaining of some outcome	Provide the elements of S
Epistemic relevance	Evidence for the obtaining of some condition	Help define P

Table 2.5

So, remember that our definition of relevance for decisions affirms that a proposition is relevant for a decision if and only if it provides reasons for or against some option. We have seen that there are four different ways in which this can happen. But our goal is to define relevance for decision *modelling*. We need therefore to update our definition of a FC-model in light of this taxonomy:

> **FC2** The *fully-considered* model is the model that includes all and only the aspects that are *constitutively* relevant for the decision.

Note that these four types of relevance play a function in relation with the elements of a decision problem: probability, utility, and the components of the decision model. This means that the identification of these four as the only types of reasons for decision making is strictly related to this way of representing decisions. It is possible that, were one to represent decisions differently, then one could identify different types of reasons in relation to the function they play in that representation. But the broad characterisation of decision problems I draw here is standard in decision theory, and, as I am trying to expand decision theory to cover a further step of deliberation, this quadri-partition should be adequate to the scope of the paper. It is important to remember, however, that I am taking the set A of acts as given - which means I am not considering cases in which the agent finds out that new courses of action are available, or that some options are unfeasible. Clearly, there may be propositions providing reasons for these changes, and being relevant to the decision for the function they play in the identification of acts. This type of relevance should be included, once the simplification adopted here is dropped.

2.6 Constitutive relevance

The problem of rational decision modelling requires an understanding of the notion of constitutive relevance, so for the current purposes we can leave investigations of epistemic and evaluative relevance aside. In this section, I will try to sketch a formal definition of the two types of constitutive relevance: CRS and CRC.

First, we need to provide some preliminary definitions. We start from a given decision D and from its set of elementary possibilities, Ω.

> *Refinement.* We can define a relation of *refinement* on the set of all possible partitions of Ω, $\Pi(\Omega)$. Given two partitions K e J, K is *at least as fine-grained as* J, $K \sqsubseteq J$, if and only if every element in K is a subset of some element of J. The relation $\sqsubseteq$ generates the complete

lattice $< \Pi(\Omega), \sqsubseteq >$, where the *supremum* is Ω and the *infimum* is the set of the singletons of Ω.

Meet. Given two partitions K and J in the lattice $< \Pi(\Omega), \sqsubseteq >$, we can identify the partition Z which is their *meet*: $Z = K \sqcap J$. This is the partition in $< \Pi(\Omega), \sqsubseteq >$ such that $Z \sqsubseteq K$, $Z \sqsubseteq J$, and for every other partition W in $< \Pi(\Omega), \sqsubseteq >$ such that $W \sqsubseteq K$ and $W \sqsubseteq J$, $W \sqsubseteq Z$.

Now we have all the elements necessary to construct the definition of constitutive relevance and to identify the criteria for the inclusion in a fully-considered model.

Let me start with CRC. As we have seen, the set A of acts is a partition of Ω, and the same can be said of the set S of states: this should be a set of mutually exclusive and jointly exhaustive alternative ways in which the world could turn out. No possibility is left out, and no two possibilities can happen at the same time.

The set C of consequences is the *meet* of these two sets: $C = A \sqcap S$. This means that C is a partition of Ω as well. This also means that C is uniquely identified by A and S: once these are provided, C is obtained with no extra input. Thus, there is no need for a notion of relevance specific to the identification of the components of C: the set that includes all and only the relevant consequences is the *meet* partition of the set A of the decision and the set that includes all and only the relevant states. If this is so, then there is only one notion of constitutive relevance – namely, the one concerning the formation of set of states. CRS will thus be referred to simply as CR.

Of course, there is a sense in which the proposition "I arrive late" is not identical with the proposition "I take the bus and there is a traffic jam", which are the act and state generating the cell of C identified by "I arrive late". However, as long as it is assumed that the two propositions are true of exactly the same set of possible worlds – that is, all the worlds in that cell – and we assume that propositions are sets of worlds, then the two are indeed the same proposition, identical to any other proposition that is true of exactly that set of worlds. Given that both probability and utility are defined over Ω and therefore are assigned to (sets of) possible worlds, then the two propositions are even practically indistinguishable, at least from the decision-theoretical perspective. One who wanted to distinguish the utility of "I arrive late" from that of the propositions "I stay longer on the bus" or "I have more time to read on the commute", which are true of the exact same worlds but toward which we can assume the agent to have a different attitude, should look for an alternative formal scaffolding.

We can thus move to the definition of constitutive relevance, CR. As partitions, the three sets of the model – A, S, C – are *meets* of bipartitions: thus, they can be modified not by adding an individual proposition, but by being refined with its corresponding bipartition. This means that whatever makes a proposition constitutively relevant is a property of its bipartition, and it is this property that we will define.

Let us begin from the set of elementary possibilities, Ω. Ω can be the *minimal* set of states: if $S = \Omega$, then its *meet* with A generates only one possible consequence for each action: each alternative action leads inevitably to one outcome, regardless of how the world turns out to be. If, on the other hand, there are some possible worlds in which some action would lead to a different outcome, then those worlds should be separated from the other ones via a refinement of Ω. This changes the composition of C, and thus potentially also the utility function defined over it. The bipartitions that are *constitutively relevant* to the decision are those the inclusion of which changes the expected utility of at least one of the actions:

> **CR.** A bipartition P is *constitutively relevant* (CR) with respect to a decision D if and only if there is an act $a \in A$ such that the difference in expected utility between a and $\bar{a}$ changes if it is calculated with Ω or with $\Omega \sqcap P$ as the set of states:

$$eu_{\Omega}(a) - eu_{\Omega}(\bar{a}) \neq eu_{\Omega \sqcap P}(a) - eu_{\Omega \sqcap P}(\bar{a}) \tag{2.1}$$

The absolute definition of CR expressed with (1) can be made comparative as well. In the equation, we have two quantities: $eu_{\Omega}(a) - eu_{\Omega}(\bar{a})$ and $eu_P(a) - eu_P(\bar{a})$. The higher the absolute value of the difference between these two quantities, the more relevant the bipartition, as this difference represents the impact that its elements have on the decision.

Of course, given our discussion of relevance we want to include in our set of states *all* the distinctions that matter for the agent's evaluation of the possible consequences, and nothing else. Thus, the fully-considered set of states is the *meet* of Ω with all the constitutively relevant bipartitions:

> **Fully-considered set of states.** A set of states S is the fully-considered set of states, S_{FC}, if and only if:
>
> $$S = \bigsqcap_{P \in \Pi_2(\Omega)} \Omega \sqcap CR(P) \tag{2.2}$$
>
> Alternatively, we can say that a set of states S is the fully-considered set of states, S_{FC}, if and only if:
>
> (i) $\forall P \in \Pi_2(\Omega)$: P is CR $\rightarrow S \sqsubseteq P$,
>
> (ii) $\nexists S' \in \Pi(\Omega)$: $S \sqsubset S'$ and $eu_S(a) - eu_S(\bar{a}) = eu_{S'}(a) - eu_{S'}(\bar{a})$.

The two definitions are equivalent: being a *meet* of partitions means being the *greatest lower bound* of those partitions, i.e. their coarsest refinement. Thus, (2) says that S is fully-considered if and only if it is the greatest lower bound of all the constitutively relevant bipartitions; equivalently, (i) says that S is a lower bound, and (ii) that it is the greatest one.

It is possible that there are no relevant bipartitions, and that the consequences of the actions are the same no matter the state of the world: then, both definitions identify Ω as the fully-considered set of states. In the first one, there would be no P for Ω to meet with; and in the second one, while for any $S, S' \in \Pi(\Omega)$ it would be the case that $eu_S(a) - eu_S(\bar{a}) = eu_{S'}(a) - eu_{S'}(\bar{a})$, Ω would be the least refined one and therefore the only one satisfying (ii).

The uniqueness and existence of S_{FC} are both guaranteed. Let us assume that S and S' are both *meets* of P and Q; then, by the definition of *meet* given above, $S \sqsubseteq S'$ and $S' \sqsubseteq S$, and therefore $S = S'$. Furthermore, given that $< \Pi(\Omega), \sqsubseteq >$ is a complete lattice, then every subset of $\Pi(\Omega)$ has a *meet*: in the absence of coarser *meet* partitions, then the set of singletons at the bottom of the lattice will be the greatest lower bound.

Some further considerations may be helpful to justify the condition at (1) as the key property to identify constitutive relevance. States are relevant to my decision insofar as they have an impact on the utility of performing some action. For instance, we can assume that the utility of carrying an umbrella is lower than the utility of not carrying it, but if it rains then the situation is reverse. In this case, the possibility of rain is relevant with respect to carrying the umbrella. On the other hand, the difference between the utility of carrying the umbrella and the utility of not carrying it is the same whether you stumble on your way out or not – stumbling may be irritating in itself, which means that the utilities are probably lower, but the difference in utility is likely to be the same regardless of whether you take the umbrella. In this case, the possibility of stumbling on your way out is irrelevant with respect to carrying the umbrella, and should not be included in the formation of the set of relevant states. Thus, a bipartition is constitutively relevant insofar as its inclusion in the construction of the model generates a set of consequences the utility of which is different from that of the consequences generated without it.

Let us have a closer look at our starting example and see how this idea works in practice. Table 1 modelled the situation of deciding whether to walk or take the bus, depending on traffic conditions: our set of states is T={traffic; no traffic}.

	Traffic	No traffic
Walk	*Arrive right on time*	*Arrive right on time*
Take the bus	*Arrive late*	*Arrive early*

If one were to assign some illustrative utilities or probabilities to the consequences and states respectively, one could end up with a matrix like this one:

	0.7	0.3
w	2	2
b	*0*	*3*

Table 2.6

Assuming that our agent is lazy and has a slight preference for taking the bus over walking, independently of any other consequence, then the matrix using Ω as our set of states could be this one:

	Ω
w	0
b	1

Table 2.7

Table 7 is our starting point, against which to test the relevance of bipartitions. In Table 7, the difference between the expected utilities of w and b is 1. Table 8 results from the *meet* of Ω and T:

	0.7	0.3
w	2	2
b	1	4

Table 2.8

Here, the expected utility of walking is $eu_{T \sqcap \Omega}(w) = 2$, the expected utility of taking the bus is $eu_{T \sqcap \Omega}(b) = 1,9$, and their difference is 0,1. Given that 0,1 is not 1, the bipartition T is constitutively relevant.

What about the possibility of rain? Let us assume that the agent assigns the following utilities to the consequences obtained under $R =$ {it rains; it is sunny}, the two elements of which are equi-probable:

	0.5	0.5
w	-2	2
b	1	0

Table 2.9

Summing the utilities of the consequences of b, the meet of this bipartition with Ω would result in the following matrix:

	0.5	0.5
w	-2	2
b	2	1

Table 2.10

Again, the bipartition is relevant: the expected utility of walking is $eu_{\Omega \sqcap R}(w) = 0$, while that of taking the bus is $eu_{\Omega \sqcap R}(b) = 1,5$, and their difference is 1,5. Given these utility assigments, both T and R are therefore constitutively relevant.

On the other hand, the possibility of finding some change in your pocket is not relevant. Let us assume that the utility of finding $5 is 5 (and 0 otherwise) and its probability 0.1, and let us look at table 9, which is the numeric *meet* of $M = \{\text{money; no money}\}$ and Ω:

	0.1	0.9
w	5	0
b	6	1

Table 2.11

Here, $eu_{\Omega \sqcap R}(w) = 0,5$ and $eu_{\Omega \sqcap R}(w) = 1,5$: their difference is thus the same as under Ω, and M should not be used to refine the set of states.

One could also wonder whether some given set of states should be coarsened. Let us take the set obtained via the *meet* of Ω, T, and M. It would yield the following matrix:

	0.07	0.63	0.03	0.27
w	7	2	7	2
b	6	1	9	4

Table 2.12

In this case, $eu_{\Omega \sqcap T \sqcap M}(w) = 2,5$, and $eu_{\Omega \sqcap T \sqcap M}(b) = 2,4$: the difference between the two is thus 0,1. However, there is a set of states that is coarser than this one, while yielding the same expected utility difference: $\Omega \sqcap T$. This means that $\Omega \sqcap T \sqcap M$ does not satisfy condition (ii), and therefore should be rejected in favour of its coarser alternative, which excludes the irrelevant bipartition M.

The ones employed in the examples are, of course, arbitrary numbers for an arbitrary utility function; they show that, for someone with that utility function, R is a relevant bipartition and M is not. This conclusion is not generalizable – it does not mean that the possibility of rain is always relevant to the decision of how to get to work. It is possible that for some agent it is irrelevant, and their utilities will reflect that.

One may also wonder whether some of the states in T should be further refined, because the actual way in which they are realised matters to the agent. For instance, the difference between heavy or light traffic may be relevant to the agent's decision. In this case, the proposition under consideration ("there is heavy traffic" or, equivalently, "there is light traffic") identifies a subset of some previous state ("there is traffic"), and the *meet* between its corresponding bipartition and T results in a split only of the "there is traffic" state. If the split results in two columns with the same expected utilities for all the actions, then the bipartition is not relevant.

A final case can be illustrated with the following example. Suppose that there is an urn with 10 red balls, 10 black balls, and 80 yellow balls. The utility of drawing a red ball is 3, while the utility of a black ball is -1 and that of drawing yellow is 0. The agent has to decide whether to play the game and draw a ball from the urn, or avoid playing and end up with a sure utility of 0 (table 10):

	0.1	0.1	0.8
	Red	Black	Yellow
Play	3	-1	0
Not play	0	0	0

Table 2.13

Is the possibility of drawing a yellow ball relevant to the decision to play? If we eliminated that possibility and redistribute the probability accordingly to the other two cases, we would obtain the following matrix (table 11):

	0.5	0.5
	Red	Black
Play	3	-1
Not play	0	0

Table 2.14

In both cases, the expected utility of playing and not playing stays the same. This is an instance of the Cancellation axiom, a principle of rational choice that has been captured by different formal notions (e.g. the extended Sure Thing principle in Savage (1954)): any state of the world under which all choices yield the same outcome should be eliminated (Tversky and Kahneman 1989). However, that does not mean that the possibility of playing yellow is irrelevant. By changing the content of the urn, one has changed the set of atoms, removing some possible worlds from Ω. Given that the composition of Ω is included in how we have defined a decision, changing Ω means dealing with a different decision, and our notion of relevance does not apply anymore.

2.7 Conclusions

The problem of providing an account of rational decision modelling is complex. The situation where an agent faces a decision can be modelled in a variety of ways and with different grains of specification. This paper has tried to provide some understanding of what I have called the fully-considered model of a decision, which is only one of the ways to coarsen the grand world model. I have argued that the FC-model is the most rational to use, at least according to an instrumental notion of rationality.

I have characterised this model as the one that includes all and only the aspects of the world that are relevant for the decision. Relevance for decisions has been interpreted as a matter of providing reasons for or against some option. This strengthens the connection between reasons and deliberation, which has largely been ignored. In presenting their decision theories, both Dietrich and List (2013) and Sher (2019) provide examples of theoretical debates in which a reason-based account of decision making could be useful. However, neither provides a clear characterisation of reasons. I have identified four different types of reasons, corresponding to four types of relevance: a proposition can be relevant to the agent's beliefs, to the evaluation of the outcomes, or as an element in the model. Propositions providing reasons of this last type are those that should be included in the FC-model, and they have been given a more precise definition.

Defining the FC-model is only one step into a complete account of decision modelling as a reasoning process. As we have seen, it is unrealistic to expect real agents to work with FC-models. More likely, they will employ other small world models. Understanding them has normative implications beyond the obvious descriptive ones. What can be said about the rationality of the models agents actually use? Describing the ideal, FC-model may help illuminate questions like this. J. M. Joyce (1999) proposes, for instance, that the agent commits to the view that a FC-model would sanction the same choice they are actually making: the only requirement of rationality in constructing the model is that the agent be justified in believing that further deliberation would not change the decision. However, what such a justification could look like is not specified by Joyce. Understanding the FC-model may help filling this gap.

Expanding from the scope of this paper, defining relevance in terms of reasons allows some further connections. One is, of course, with the literature on practical reasons. An important distinction that is often drawn in there is one between explanatory and normative reasons (Dancy 2000a). While decision models can be constructed with either type, the inclusion of normative rather than explanatory reasons may lead to different recommendations. Different notions of rationality and different philosophical enterprises may be interested in adopting either type.

Moreover, introducing reasons points to notions of rationality beyond the purely instrumental one. Gilboa and Schmeidler (2001) define rationality as a robustness of some sort: according to them, a decision is rational if the agent does not wish to change it upon analysis. Something similar is formalised by Grant and Quiggin (2008) as the requirement of sequential consistency, which is respected whenever the agent still regards a past decision as ex-ante optimal in light of increased awareness. Finally, there is a trend (Dennett 1987; McGeer 2007) that sees rationality as rationalisability, and norms of rationality as norms to make sens of each other's behaviour. In all these cases, rationality requires to provide reasons for our decisions. All this is material for further work.

Chapter 3

Non-cognitive uncertainties in evidence-based decisions

Abstract – The increasing success of the evidence-based policy movement is raising the demand for empirically informed decision making. As arguably any policy decision happens under conditions of uncertainty, following our best available evidence to reduce the uncertainty seems a requirement of good decision making. However, not all the uncertainty faced by decision makers can be resolved by evidence. In this paper, we build on a philosophical analysis of uncertainty to identify the boundaries of scientific advice in policy decision making. We argue that the authority of scientific advisors is limited to cognitive uncertainty and cannot extend beyond it. While the appeal of evidence-based policy rests on a view of scientific advice as limited to cognitive uncertainty, in practice there is a risk of over-reliance on experts beyond the legitimate scope of their authority. We conclude by applying our framework to a real-world case of evidence based policy, where experts have overstepped their boundaries by ignoring non-cognitive types of uncertainty.[1]

[1] A different version of this chapter has been co-authored with Mattia Andreoletti and published on *Perspectives on Science*, *30*(2), 305-320.

3.1 Evidence-based decision making

Uncertainty is a crucial feature of decision making. Decisions are aimed at reaching or avoiding some result, but we may be uncertain about which of the actions at our disposal will bring about the desired outcome. We typically ignore the exact consequences of our actions, just as we cannot be sure about the circumstances in which they will happen. Indeed, arguably any decision we face happens under some degree of uncertainty. If this is so, then it seems that good decision making should try to reduce that uncertainty, and make choices that are as informed as possible.

With regard to policy decision making, A. V. Banerjee and Duflo (2011) claim that many policies fail due to ignorance – specifically, ignorance about what actually works. Under this view, policy makers are at fault in their decision making process because they do not strive to reduce the uncertainty surrounding their decisions. Instead, to make good policy decisions they should listen to scientists, who may be in the position to provide evidence in support or against the options at stake. The view that policy decisions should be based on scientific evidence has fuelled a movement that has enjoyed a rising success in the last few decades (A. V. Banerjee, Duflo, and Kremer 2016; Ravallion 2020), culminating in the recent Nobel Prize (2019) awarded to some of its most prominent figures.

The advocates of this account, known as Evidence-Based Policy movement (EBP), call for the incorporation of rigorous and robust scientific evidence in policy decision making. Its success has led to the proliferation of empirical studies in social policies, and to the increasing authority of experts (economists) in policy decisions. This trend has reignited the debate over the relationship between science and policy in general, and between evidence and values in particular. While the traditional debate focused on the role of values in evidence gathering (e.g. Douglas (2000), Longino (1990), Rudner (1953), and Steel (2010)), the rise of EBP has fostered new sets of questions. Some authors have questioned the ethical validity of randomisation (Ravallion 2014, 2020), while others have illuminated the relation between policy evaluations and the empirical methods used (Dede 2019). Finally, some have investigated the trade-offs between epistemic and non-epistemic values in evidence for policy (Khosrowi 2019; Khosrowi and Reiss 2019).

Here, we want to step back from the debate over empirical methodology and put the spotlight on the role scientific advice has in making evidence-based policy decisions. Policies are not decided by the scientists constructing the evidence. However, if good decision making should try to reduce the uncertainty surrounding the decision, then decision makers should listen to what science has to say on the matter at stake. But what does listening to science mean?

Scientific advice could enter the decision making process as a recommendation about what should (not) be done. In light of their knowledge, experts could have reasons to support or oppose some courses of actions, and policymakers could listen to science in the sense of making decisions according to what the scientific community suggests. A second way in which scientific advice could enter the decision making process is not as opinion about the decision itself, but as information concerning some aspects of the decision. In this second sense, policymakers would listen to science insofar as the information they use in their deliberation comes from solid evidence.

In this paper, we build on a philosophical analysis of the uncertainty faced by policy decision makers to understand the boundaries of scientific authority in decision making. While EBP may ground its appeal on an (often implicit) assumption of the role of scientific advisors as neutral providers of evidence, their practice risks to be closer to the first role described. We start by introducing the concept of non-cognitive uncertainty, and we focus on two types of non-cognitive uncertainty that play a crucial role in decision making and that are particularly relevant in policy contexts. The presence of these uncertainties set the boundaries for the authority of scientific advisors in policy decision making, which is limited to a types of uncertainty. We then proceed to discuss some limitations in the practice of EBP that make evidence-based decision making more exposed to the first type of scientific advising. We conclude by discussing a real-world case in which EBP practitioners have overstepped the boundaries set by non-empirical uncertainties.

3.2 Non-cognitive uncertainties

Decision making is a process of resolving uncertainty. Trivially, the agent facing a decision does not know what to do, and is therefore in a situation of practical uncertainty (Peter 2021), i.e. uncertainty concerning the course of action to take: making a decision is resolving practical uncertainty. Understanding the nature of this uncertainty is therefore crucial for understanding good decision making.

Since its origins (Ramsey (1926) and Savage (1954)), contemporary decision theory has been developing probabilistic tools to tackle uncertainty. The standard view is that the uncertain aspect of decisions can be entirely captured by one probability function over possible states of the world. Even though in most real-life cases it is impossible to assign precise probabilities to alternative states, sophistications of standard decision theory aiming to capture this more severe uncertainty still strive to do so in probabilistic terms, e.g. using families of probability functions (e.g. J. M. Joyce (2010) and Levi (1983)) or second-order weights on probabilities (e.g. Chateauneuf and Faro (2009), Gärdenfors and Sahlin (1983), and Klibanoff, Marinacci, and Sujoy Mukerji (2005)).

The assumption behind this enterprise is that the only obstacle on the decision maker's path is the lack of adequate information – albeit with different degrees of severity. If all the indecision were due to uncertainty about what is or will be the case, then it would be entirely resolved once the adequate empirical knowledge were provided. In principle, there is some evidence that could effectively resolve the agent's practical uncertainty, provided that the relevant proposition is epistemically accessible.

Then, the best option would be the one leading straightforwardly to the preferred outcome.

Unfortunately, practical uncertainty is not entirely reducible to lack of information. In itself, knowledge of the state of the world does not imply any choice: the agent may still be unsure about what to do even under complete certainty. Decision theory focuses on the part of practical uncertainty concerning the agent's descriptive judgements, i.e. their beliefs about what is or will be the case (Bradley and Drechsler 2014). We will call uncertainty concerning attitudes, like belief, that purport to represent reality *cognitive*. However, there is growing work showing that not all the uncertainty is of this kind. For instance, in his typology of uncertainty, Hansson (1996) identifies four components of what he labels "great" uncertainty, not all of which concern beliefs about the state of the world. Similarly, Helgeson (2020) investigates "deep" uncertainty in the structuring of decisions. Bradley and Drechsler (2014) identify uncertainties concerning normative and modal judgements, as well as descriptive (see also Dietrich and Jabarian (2018)). Non-cognitive uncertainties are those types of uncertainty that do not concern what the actual state of the world is, and that therefore cannot be solved by evidence.

If there are non-cognitive types of uncertainty, then supporting a policy "because it is evidence-based" may be misleading. Support for a policy means choosing that policy over some alternatives: it implies a decision. But reaching a decision means that all the uncertainty around it has been accomodated, one way or another. If, as we say, some of this uncertainty is insensitive to evidence because it is non-cognitive, then no policy can ever be chosen entirely on evidential grounds. Choosing a policy will always require some non-cognitive judgements, no matter the amount of evidence available.

In the next two sections, we focus on two types of non-cognitive uncertainty that play an important role in decisions for policy making, evaluative and model uncertainty. In each case, we present the uncertainty, we clarify the sense in which it is non-cognitive, and we highlight the aspects that are relevant to policy making.

3.2.1 Evaluative uncertainty

The agent picks a certain option when they find it more desirable than the alternatives. If they are uncertain with respect to the desirability of some alternative, they are in a situation of evaluative uncertainty. According to Bradley & Drechsler, this uncertainty (which they label *ethical*) arises when the values used to assess the desirability of the different alternatives are unknown or non-existent (2014: 1237). The extent to which this uncertainty is non-cognitive depends on one's general view on evaluative judgements.

Ethical subjectivism, which is the most prominent position within decision theory (Bradley and Drechsler 2014), takes evaluative uncertainty to concern one's own judgements of desirability. Given that these judgements do not purport to represent reality, this position makes evaluative uncertainty a non-cognitive uncertainty. As long as the evaluation of consequences is based entirely on the agent's tastes, then subjectivism makes this uncertainty quite unusual, as it would require one to be uncertain about one's own tastes. But if the evaluation requires normative (aesthetic, moral) considerations, then the agent may be uncertain about which considerations to adopt and which values these assign. Evaluative uncertainty would then be both plausible and non-cognitive, as long as one is a non-cognitivist about normative judgements. If the non-cognitivist is right, then there is no procedure-independent correct answer to the question of how to evaluate the available alternatives. As there is nothing to discover about the desirability of the options, then no amount of evidence could solve the uncertainty. Once all the relevant information is available, whether something is more valuable than something else remains a non-empirical question.

On the other hand, one can be a cognitivist and take evaluative uncertainty to concern beliefs over some objective normative facts. Since it is a matter of truth whether something is good (tasty, beautiful...) or not, then desirability can be the object of uncertain beliefs. This view reduces evaluative uncertainty to a standard cognitive uncertainty about a specific class of beliefs – but it does so at the cost of assuming realism about what is good. Alternatively, cognitivists assume that evaluative uncertainty concerns beliefs not over normative properties, but over the agent's tastes, which are taken to be objective facts about a person.

Without trying to settle the debate over which of these views is the adequate understanding of evaluative uncertainty, let us now explore the implications they have for policy making. To do so, we assume that the crucial difference between individual decision making and policy decision making is that the latter concerns a plurality of stakeholders.

Cognitivism about normative properties implies that there is a correct way to evaluate the different alternatives. If the decision maker is uncertain, then they can find out something more about what is valuable. In sum, they should treat their uncertainty over normative judgements just as their uncertainty over descriptive judgements: as something that can, at least in principle, go away with the right information. This is so whether the decision concerns individual choices or policy options. But this requires realism over normative properties – a position that may be at odds with the practices of democracy, which typically assume at least some pluralism.

Cognitivism about tastes has some more interesting implications. If the preferences of all the stakeholders are relevant for the decision, and not just the personal preferences of whoever happens to be making the decision for everybody, then evaluative uncertainty in policy contexts amounts to uncertainty over the stakeholders' evaluation of the alternatives. According to this view, this is cognitive uncertainty over some factual features of the people involved: each stakeholder has some tastes, and evaluative uncertainty means not knowing some of these tastes. However, while evaluative uncertainty could be reduced to standard cognitive uncertainty in single-agent contexts, in policy making it involves an additional level: even if the decision maker had all the information about everybody's tastes, they would still need to aggregate all that information into one single evaluation. But there is no independently correct procedure to do so. Non-cognitive uncertainty comes back at this higher-level decision making.

Let us now look at the fully non-cognitive view, namely ethical subjectivism. This view claims that there is no correct answer to the question of how desirable some alternative is. Therefore, it is possible that some stakeholders disagree on how to evaluate the alternatives, and that this disagreement is irreducible: it may be due to different subjective evaluations, none of which is inherently wrong. People may legitimately disagree on the evaluation even starting from the same background knowledge.

In this case, the answer to the question of desirability must come from a certain side of the disagreement. It is a situated answer, coming from a specific standpoint on the issue. On whose evaluations should the decision be based? Whether the evaluation is the decision maker's, the experts', the result of some averaged aggregation, or is settled via standard democratic procedures – it is still only one of the admissible (i.e. not incorrect) evaluations. In the absence of an independently correct answer, the uncertainty over the question of desirability is not sensitive to empirical evidence.

3.2.2 Model uncertainty

The model of a decision comprises of a set of available actions, a set of consequences, and a set of possible states of the world, i.e. the worldly contingencies on which the outcomes depend. Standardly, decision theory takes the model as given, and does not question the elements included in or excluded from it. Nonetheless, in real-world decisions models are not ready-made: the agent has to build their own. Model uncertainty arises whenever the decision maker is aware of the possibility that they have not included all the relevant elements in their model.

The possibility of uncertainty over the relevant states has given rise to a prolific literature on decision making under unawareness, i.e. in those circumstances where the agent is aware that they may be unaware of some contingency (Karni and Vierø 2017; Schipper 2014). In this basic form, this type of uncertainty seems to be cognitive, albeit of extreme severity: it concerns the agent's beliefs over the world. Far from being able to assign precise probabilities, the agent has no information whatsoever regarding those contingencies of which they are unaware. However, it is a matter of fact whether those contingencies obtain or not, and in principle there is some evidence that could resolve their uncertainty.

There is, however, an important non-cognitive element. It is clear that the agent cannot be aware of all possible contingencies, let alone include them all in their state set. The question that the agent faces when constructing their model is not whether they have included all the possible things that may be the case or that may happen – rather, they face the question of what possibilities to include and what to exclude from their sets. This means that something relevant may be left out of the model not only because the agent is unaware of it, but also because the agent excluded it overlooking it as something irrelevant to the matter at hand.

Since it is impossible to include every eventuality in the model, then it is always possible that the agent has omitted something relevant. Thus, model uncertainty is pervasive in decision making, simply because any decision requires a selection of the relevant factors on which to base it, and the selection may turn out to be inadequate. However, judgements of relevance are not entirely cognitive. While to some extent they do concern beliefs (e.g. beliefs over causal relations between actions, contingencies, and outcomes), they are also the results of subjective considerations of the agent about what matters to them: what counts as relevant to the decision ultimately depends on what the agent hopes to achieve with the decision. To the extent to which these sorts of judgements can be interpreted as involving non-cognitive attitudes, then model uncertainty is a non-cognitive type of uncertainty. If this is so – if judgements of relevance are at least to some degree non-cognitive – then the question of what to include in the model does not have a correct answer. This means that, once again, it is possible that different agents disagree on what should be included, and that this disagreement is irreducible.

The same reasoning applies to the inclusion of outcomes of the possible actions in the set of consequences. The agent's decision does not depend only on factual contingencies, but also on which consequences they expect from the performance of each alternative option. Now, actions can have a variety of effects. The agent may be unaware of some of these – they may not expect a certain option to result in a certain effect, and therefore exclude it from their considerations. But they may also exclude some effects as irrelevant for their evaluation of the options. Again, this judgement of relevance is subjective, and people can disagree on which effects matter and which do not.

Therefore, the selection of the elements – contingencies and consequences – on which to base the decision is the result of an agent-relative judgement of relevance. This has direct implications for policy decisions since, as we have noted, policy decisions impact on a plurality of stakeholders. Different actors may have different opinions on what matters in a certain decision: they may consider different consequences of the policy options as the ones that are truly at stake, and they may have different ideas as to which facts of the world these consequences depend on. As there is no single correct answer, different positions can be equally legitimate: they may represent different priorities on the matter at hand.

For these reasons, the existence of robust evidence about some of the policy effects does not imply that the evidence effectively settles the issue. Some stakeholders may think that it does, because the only thing that matters is whether a certain action produces some specific set of effects. Others may claim that it misses the point entirely, because those effects are not really what the issue is about. Their disagreement is not due to lack of sufficient information, but to different priority settings.

And priority setting has more to do with non-cognitive attitudes than with cognitive ones.

3.3 The role of scientific advice

In light of the uncertainties we have reviewed, we can say that an agent's practical uncertainty with respect to a decision comprises at least three different types of uncertainty. An agent may be uncertain about what to do because they lack some crucial information about the actual state of the world (cognitive uncertainty), because they do not know how to evaluate the possible consequences of their actions (evaluative uncertainty), or because they are not sure about which contingencies and consequences they should take into consideration (model uncertainty). In the context of policy making, the presence of a variety of stakeholders implies that both evaluative uncertainty and model uncertainty may be ascribable to some irreducible disagreement over value systems and priorities. As this disagreement is not over some matter of fact, for which there is a correct solution, it cannot be settled by evidence.

We are now in a position to better qualify the role of scientific advisors in the reduction of practical uncertainty. We have seen that there are two possible ways in which their role in decision making can be cashed out: they can provide opinions as to which option should be pursued or avoided, or they can abstain from opinions and limit themselves to the provision of data and information. The scientific advisor who followed the first route would need to have an opinion on the decision at stake. In order to form that opinion, they will have to solve the practical uncertainty surrounding the decisions: to say that option *a* is better than option *b* they have to solve the decision problem. In order to do that, the scientific advisor has to face both cognitive and non-cognitive uncertainties. On the other hand, scientific advice along the second route does not require the resolution of practical uncertainty, and thus of non-cognitive uncertainties.

As any stakeholder, the scientific advisor can form an opinion on the decision at stake. In order to do so, they will rely on their expertise to resolve the cognitive components of the uncertainty. But as for the non-cognitive components, they are in no special position with respect to other stakeholders. The specific authority of the scientific advisor is epistemic in nature: it is authority over the formation of beliefs about the state of the world. Epistemic authority does not automatically grant practical authority, i.e. authority over what should be done. Thus, the opinion of the scientific advisor formed on their scientific knowledge is not more valuable than that of other stakeholders with the same scientific knowledge at their disposal. If this is so, then the role of scientific advisory in evidence-based decision making is to provide the scientific knowledge with which to reduce the cognitive uncertainty, rather than fully formed opinions on the decision itself.

The appeal of EBP presupposes this second account of scientific advice. In this account, scientific advisors do not address questions regarding the desirability of policy interventions or their different priority. This allows EBP advocates to claim that their judgements regarding policy interventions are essentially neutral. Experts' judgements are authoritative insofar as they concern the area on which experts hold justified authority. Insofar as these judgements are only relative to cognitive uncertainty – that is, insofar as they are provisions of evidence and data – then they are authoritative, because experts hold justified epistemic authority on empirical issues.

However, this view seems to be in contrast with what actually happens in contemporary scientific advising. Nowadays, science is pervasive in every aspect of our society, and scientists occupy influential and important roles in governments. This creates political pressure on the decision makers. Even if scientific advisors limited themselves to "neutral" (i.e. limited to cognitive uncertainty) advising, it might be "uncomfortable and politically risky" to ignore it (Douglas 2009, p. 43). If the advice seems to favour some option, then in the eyes of the general public decision makers should privilege that option as this is "what science says". So, even if scientific advisors themselves only provided evidence, their advice may be transformed into a fully formed opinion, and decision makers are likely to be pressured into following that opinion.

Moreover, there are many cases where clear recommendations on what to do is openly requested to scientific experts. This is the case, for instance, of regulatory decisions on new pharmacological treatments. Governmental agencies such as the US Food and Drug Administration (FDA) and the European Medicines Agency (EMA) are political institutions which nonetheless grant market approval decisions on scientific recommendation. For each request of drug approval, the regulatory agencies summon advisory committees of experts to assess "whether the safety and effectiveness information submitted for a new drug is adequate for marketing approval" (Andreoletti and Teira 2016). The committees are asked to vote for or against. Although their conclusions are not mandatory for the agencies, they are almost always accepted. According to empirical studies of scientific advisory and regulatory decisions, the agencies' decisions are very consistent with the advisory committee votes (see e.g. Zuckerman (2006)).

Finally, the rhetoric around EBP is itself ambivalent. On one hand, it criticises policy making as being ideologically driven and biased by preconceptions, stressing the need of the neutral eye of science to identify effective policies. In doing so, it seems to blame policy making for introducing ideology into the realm of cognitive uncertainty, where in principle objectively correct answers are possible. On the other hand, it conveys the idea that we should make policies that the evidence shows to work. But that a policy works is an incomplete description: a policy works on some aspect, it works in some way and for certain stakeholders. Overlooking these specifications assumes that however they are resolved in the collection of evidence is what should matter in the policy context. In doing so, EBP proponents are not just providing useful data to make an informed decision, but they are effectively shaping the choice. The rhetoric of "what works" tries to cleanse policy making from ideology beyond cognitive uncertainty, where value-laden stances are unavoidable.

The focus on effectiveness in contrast with ideology may lead to considering it as the main criterion to choose a policy. In this case, scientific advisors may end up transforming their judgements of effectiveness into opinions of which policy should be implemented, thus crossing the boundaries of cognitive uncertainty. The experimental approach itself can then start to be seen as an instrument of persuasion, to push policy makers in one direction thanks to the epistemic authority of scientific methods:

> One way to interpret the series of (...) studies is as a process of persuasion at scale: the experimental approach played not only an evaluation role but also an instrumental role in fostering acceptance of the policy by the government. (...) From that perspective, the experimental approach is a little like opening a jammed door with a pry-bar. First you stick the bar in a little crack, and get a little traction. Then you move to another location, and get a little more traction. When you've got a little more purchase, you can jam in a bigger pry-bar and really tug hard. (...) At some point, the leverage is great enough that you can throw the door open. Sequential experimentation becomes a political economy tool for getting momentum for policy change. (A. V. Banerjee, Banerji, Berry, Duflo, Kannan, Shobhini Mukerji, et al. 2017, p. 31).

The scientific advisor has decided that the policy they are evaluating is one that should be implemented, and uses the experiments on which this judgement is based to promote it. Evidence can thus become a political instrument to support different policies, even though without some (value-laden) evaluation of its consequences and some (value-laden) assessment of its relevance it cannot directly support anything. In the next section, we present a case in which EBP practitioners moved from the role of providers of evidence to policy supporters, thus moving beyond cognitive uncertainty.

3.4 A real world application: the case of TaRL Africa

So far, we have analysed two types of non-cognitive uncertainty and the boundaries they set for the role of scientific advice in policy making. It is now time to see how this theoretical framework can help highlight how EBP can overstep these boundaries by looking at a real case of evidence-based policy making.

A prominent promoter of the EBP movement is the Abdul Latif Jameel Poverty Action Lab (J-PAL). Founded in 2003 by MIT professors Abhijit Banerjee, Esther Duflo, and Sendhil Mullainathan "with the goal of transforming how the world approaches the challenges of global poverty" (J-PAL 2022), it has grown to involve over 400 professionals. Nowadays, programmes based on its evaluations have reached more than 400 million people around the globe. One of J-PAL's flagship works is "Teaching at the Right Level" (TaRL) [2], an education programme developed in collaboration with the Indian NGO Pratham.

In the early 2000s, Pratham started to experiment a new pedagogy to fight low literacy among children (A. V. Banerjee, Banerji, Berry, Duflo, Kannan, Shobhini Mukerji, et al. 2017). The core idea was to divide the kids, for some part of the day or for some period of the year, according to their abilities rather than their age: this was supposed to facilitate the children's literacy acquisition. As the programme started expanding, researchers from J-PAL got involved with Pratham to evaluate the effectiveness of their method and guide its scaling-up. Several years and six randomised experiments later (A. V. Banerjee, Banerji, Berry, Duflo, Kannan, Mukherji, et al. 2016; A. V. Banerjee, Banerji, Duflo, et al. 2010;

[2]https://www.teachingattherightlevel.org/

A. V. Banerjee, Cole, et al. 2007), TaRL reached a national scale, involving millions of kids.

In Delhi, government schools applied Pratham's methods in 2016. But TaRL was seen as implementing a policy of segregation in classrooms, leading some parents to file a lawsuit against Delhi government (Singh 2019). In fact, teachers testified that level-based division has an impact on the identity formation of children, that were now often identified with their level. According to them, this translated into bullyism and demotivation. Moreover, teachers accused that, within this discriminatory system, students classified as "bad learners" were given less interesting and engaging teaching activities (Kalra 2019).

A. V. Banerjee, Banerji, Berry, Duflo, Kannan, Shobhini Mukerji, et al. (2017) present the series of experiments conducted with Pratham precisely as a virtuous example of how to successfully scale up an intervention, from piloting tests to policy implementations at the large scale. The basic gist of Pratham's TaRL pedagogy was tested in a variety of contexts, proving effective in Indian states with very different socio-economic profiles. For them, the challenge was to find ways to implement TaRL core principles outside of the specificities of Pratham's own programme. The evaluations showed which implementations worked better. Furthermore, they showed that problems arising when the programme was implemented in government schools could be overcome with constant support from Pratham's staff.

In virtue of the results provided by the evaluations, their authors provide full support to policies based on TaRL. Recently, J-PAL and Pratham started the joint venture "TaRL Africa" to export Pratham's pedagogy to Africa. On January 15th, 2019, the philanthropic collaborative Co-Impact granted a millionaire commitment in support of TaRL Africa, which was chosen with other four initiatives from a pool of 250 candidates. Iqbal Dhaliwal, Executive Director of J-PAL, commented on the commitment as a victory for evidence-based policy making:

> This grant represents the critical importance of using evidence from rigorous impact evaluations to drive decision making. (Pratham 2019)

However, the experiments provide reasons for a policy to the extent that policy choice requires the reduction of cognitive uncertainty. But, as we have argued, there are uncertainties that cannot be resolved empirically.

TaRL Africa is a case of policy promotion, rather than just intervention evaluation. As we have seen, winning the Co-Impact funds was done at the expense of other policy proposals. The researchers are supporting a policy over other policies, addressing different aspects of the scholastic context or different social issues altogether: their support is the result of a decision. As such, it has involved judgements of relevance and desirability, as well as of effectiveness. Let us look more closely at how non-cognitive uncertainties are at play in the TaRL case.

As we have seen, evaluative uncertainty concerns the evaluation of the outcomes of the intervention. We can imagine, for instance, that the teachers or the families that will be touched by TaRL Africa may not evaluate it positively, because in their view bullyism is more worrying than slow literacy acquisition. If there is disagreement among the stakeholders as to the value of the outcomes of the intervention, then its promotion implies the adoption of only one of the available evaluations. But since the disagreement may be due to different priority settings, then the choice of an evaluation may be the choice of a priority setting. And that is not an empirical issue. Thus, the promotion of a policy based on an intervention that raised mixed feelings requires a normative choice over what to prioritise in the intervention outcomes. Indeed, if ethical subjectivism is right, the promotion of any policy comes from a specific standpoint.

As for model uncertainty, it comprises two levels. On the higher level, Co-Impact is not limited to education policies. Thus, promoting a policy targeting education over policies targeting other social issues implies that education is more relevant. On the lower level, education itself is made of a variety of components. TaRL has mixed effects: while it has positive effects on the rate of literacy acquisition, it has been claimed to have negative effects on children's self-esteem and on bullyism. The promotion of TaRL in the name of the former set of effects implies that they are more relevant than the latter. But this is not an empirical judgement, and as such it cannot be settled by evidence. The promotion of a policy based on TaRL is therefore the result of a normative choice as to what matters in education, as well as to what matters in development. In this case, choosing TaRL because experiments show its effectiveness in literacy acquisition means that literacy acquisition is a more relevant effect than others. But since judgements of relevance are at least partly non-cognitive, then disagreements are both possible and legitimate, and cannot be settled by resorting to evidence.

3.5 Conclusions

Decision makers often lack complete information on the key contingencies of their decisions. Systems on which policies intervene are complex and the spectrum of choices available is very wide: most of the time then decision makers get hung up on the nitty-gritties of a conundrum, and resorting to scientific evidence has been offered as a way out of this impasse. The surge of the Evidence-Based Policy movement in the last decades has sparked academic and non-academic debates on the legitimacy of its approach. In this paper, we looked at the role of scientific advice in evidence-based decision making. We framed decision making as a process of resolution of practical uncertainty, and showed that this comprises both cognitive and non-cognitive uncertainty. We used this framework to claim that, while the appeal of EBP rests on an understanding of scientific advice as limited to the reduction of cognitive uncertainty, in practice its impact on decision making tends to be larger. The case of TaRL Africa shows how EBP researchers may move beyond cognitive uncertainty to promote a policy that their evidence proves to be effective. As promoting a policy implies having chosen that policy, it faces evaluative and model uncertainty. Their resolution requires judgements of relevance and desirability that are outside of the epistemic authority of scientific advisors.

Chapter 4

COVID-19 and the need for inclusive policy-making

Abstract – The COVID-19 pandemic has presented the world with a series of new challenges, but the policy response may be difficult due to the severe uncertainty of our circumstances. While pressure to take timely action may push towards less inclusive decision procedures, in this paper I argue that precisely this uncertainty provides both democratic and epistemic reasons to include stakeholders in our collective decision making.[1]

4.1 Introduction

The COVID-19 pandemic has taken the world aback. It has shaken almost every aspect of our individual and social life, presenting us with severe, novel challenges. These unprecedented circumstances are so fraught with uncertainty that making good decisions may be difficult.

[1] A version of this chapter has been published on *Erasmus Journal for Philosophy and Economics, 14*(1), 128-137.

And yet, the pandemic requires timely and apt responses. In the face of this uncertainty, how should policy makers make decisions? While the pressure for timely responses may push towards swifter and less inclusive decision procedures, I argue that our current uncertainty provides both democratic and epistemic reasons for the importance of stakeholders' inclusion in collective decision making.

Bayesian decision theory provides important tools to deal with decision making under uncertainty, but it assumes that the uncertainty faced by the agent can entirely be represented in probabilistic terms. However, the uncertainty faced by policy makers during a pandemic goes beyond this standard view: not only because it may be too severe, but also because it concerns aspects of the decision that are excluded from it. As we will see, the uncertainty concerns also the construction of decision models and the evaluation of the possible consequences of a policy – both issues over which people may have irreconcilable disagreements due to their values and priorities. For this reason, these uncertainties call for the inclusion of all stakeholders in the decision making processes. But inclusion is valuable even when facing standard uncertainty. Considerations of accuracy may point against this conclusion, privileging expertise over inclusion; however, the exclusionary nature of expertise keeps removed from our collective efforts of knowledge production precisely those communities that are more severely hit by the pandemic.

The paper is structured as follows. In the next section, I present the standard view of uncertainty as treated in decision theory, and discuss how it may be limited, while in section 4.3 and 4.4I move to model and evaluative uncertainty respectively. All these uncertainties are made particularly severe by the pandemic, especially due to the complexity of a situation that involves many interconnected aspects and may touch stakeholders differently. Finally, in section 4.5 I explore two different arguments for inclusive decision making that are suggested by the nature of the different uncertainties discussed.

4.2 Cognitive Uncertainty

When making decisions, policy makers face situations of practical uncertainty, i.e. uncertainty concerning what to do (Peter 2021). Bayesian decision theory suggests that the way to resolve this uncertainty is to choose the option from which the agent expects the best outcome, given their beliefs about the world. The value of the different outcomes is represented by the agent's utility function, while their beliefs about the possibility of the contingencies on which these outcomes depend are captured by a probability function. Then, the agent should pick the available option leading to the highest expected utility given the probability of different scenarios.

This picture assumes that the agent's uncertainty can be entirely captured by the probability function representing their beliefs over possible states, an uncertainty which I will call *cognitive*. While very elegant, this assumption is rather strict. By this, I do not mean to criticise the widespread Bayesian tenet that degrees of belief can be represented by probabilities. Rather, I will focus on the uncertainty that is left out of the picture even on the assumption that it is possible to represent degrees of belief with probabilities. The uncertainty may be too severe for the agent to be able to form consistent degrees of belief, or it may involve attitudes beyond belief: thus, if it is too severe, or if it is not of the right kind, uncertainty cannot be captured by probabilities even holding the Bayesian tenet. Let us start from the first possible limit, severity. In the next two sections, we will explore types of uncertainty that may cover values as well as beliefs.

Representing degrees of belief with probabilities requires, of course, that the agent is able to form degrees of belief. Yet, the severity of the uncertainty may be such that this is not possible. For instance, we may have rather specific degrees of belief in the effectiveness of a certain vaccine, and it may be natural to represent these with the probability rate of its success. But this is not the case for all events. The agent may have some degree of belief in the eventuality that the government will fall due to its inadequacy in organising the vaccination campaign, but be little confident of its accuracy – and surely we do not have probability measures of its likelihood. In some more abstract cases, the agent may be unable to form degrees of belief at all, for example with respect to the possibility of the insurgence of a new pandemic or to the impact that restrictive measures may have on an insurgence of nationalisms. And finally, there are also all the contingencies the agent is unaware of, i.e., those things that the agent does not even know that they do not know (Rumsfeld (2002)'s famous "unknown unknowns").

These illustrative cases show that decision making in response to the pandemic involves uncertainty that is often more severe than mere probabilistic risk. While classical expected utility theory may be inadequate to tackle decisions under more severe degrees of uncertainty, numerous sophistications have been developed to tackle these cases while still holding on the Bayesian tenet (e.g. Bradley (2017), Ghirardato (2001), and Karni and Vierø (2017)). Thus, even though it cannot be entirely reduced to probabilistic risk, severe uncertainty can still be addressed within some expanded version of Bayesian decision theory.

However, in this discussion we have assumed two conditions: first, that the decision problem is already modelled; second, that the utilities are given. In the next two sections, I will explore each of these conditions in turn, showing how they lead to further uncertainty, and how they point to stakeholder inclusion as a requirement for decision making during the pandemic.

4.3 Model Uncertainty

Formal theories require that decision problems be structured in models that they can take as input. As Bayesian decision theory calculates the best option on the basis of the utility of the outcomes of the alternative options and of the probability of the contingencies on which these outcomes depend, then it can only be applied to problems that have been structured in a set of available alternative courses of action, a set of outcomes, and a set of contingencies (or states). But real life problems do not come with pre-formed sets. Thus, the theory can only be applied once the problem has been modelled – yet the theory itself does not provide any instruction on how the model should be constructed. Insofar as agents can be uncertain about the right way to model their decision, then practical uncertainty cannot be entirely resolved within the boundaries of Bayesian decision theory. And indeed, there are many reasons to expect uncertainty over decision modelling in policy responses to the pandemic.

As the current circumstances are unprecedented in our lifetimes, policy makers are exploring options that are outside of their usual protocol: extraordinary times require out-of-the-ordinary solutions. Restrictive measures that were unheard of in democratic societies are now commonly debated and implemented; creative solutions have been devised to host events without compromising safety; many activities have had to change shape and move online. This means that policy makers work with an open set of acts: rather than sticking to the old guidebook, they may have to come up with novel solutions to novel problems. For this reason, they may be uncertain as to whether they are really considering all the alternatives at their disposal, as well as doubting whether some standard strategy is still truly available.

Even in the identification of the set of outcomes there is an uncertain aspect. Any given action leads to an indefinite number of consequences: my decision to go to a certain café will have the consequence of me drinking a certain blend of coffee —- but also of sitting on a stool, crossing the street, and entering the café with my left foot, among countless other actions. Not all these consequences will be relevant for my choice. The exact number of steps to get to the café is probably not relevant, while the overall distance may be. Some people may care about the comfort of their seat, while others may only care about the quality of the coffee. Clearly the agent should not consider all the minute consequences of an action, but only those relevant to the decision. The same applies to the description of the different possible scenarios in which the action may occur, and therefore on the identification of the set of contingencies: each state of the world is constituted by a myriad of factors, the relevance of which depends on the grain of specification one adopts and on the selection of consequences one cares about.

This means that there is no unique way to model a certain decision problem. If this is so, then agents may well be uncertain about their model, and wonder whether they have included all (and only) the relevant options, outcomes, and contingencies. I will call this uncertainty *model uncertainty*, to distinguish it from the cognitive uncertainty that we have seen addressed in Bayesian decision theory.

Model uncertainty is particularly acute under the pandemic, as policies tend to intervene on complex situations made of several interconnected aspects. A policy may impact both public health and the economy, while also having environmental and political consequences. Of course, one could strive to construct models that are as comprehensive as possible, but at some point a selection is inevitable – and the more complex the situation, the more controversial the selection may be, given that judgements over what is relevant to a decision may depend on one's values and priorities. Moreover, given the severity of our cognitive uncertainty, it is possible to end up discarding some factors that may turn out to be relevant after all.

4.4 Evaluative Uncertainty

Once the decision is modelled, the resolution of the agent's practical uncertainty requires more than just probabilities. As decisions are the result of both the beliefs and the desires of the agent, Bayesian decision theory requires a utility function to represent the agent's preferences as well. While probabilities (or lack thereof) are themselves representations of uncertainty, utilities are taken to represent value with certainty, as if desires were not subject to uncertainty the way beliefs are.

Instead, I will refer to the predicament of the agent that is not certain about the value of consequences as *evaluative uncertainty*, to distinguish it from both cognitive and model uncertainties. Evaluative uncertainty is not usually included in Bayesian decision theory. Some authors have proposed extensions of the theory with families of utility functions, but their efforts aim at tackling incompleteness of preferences rather than uncertainty (e.g. Bradley (2009), Galaabaatar and Karni (2013), and Ok, Ortoleva, and Riella (2012)). This is not surprising: it may be somewhat counter-intuitive to think of uncertain desires, as these seem to be either trivial or reducible to cognitive uncertainty. Either the agent is unsure about their own tastes, in which case the uncertainty should be easily remedied with some more introspection; or they are uncertain about some factual properties about themselves or about some normative facts, in which case evaluative uncertainty can ultimately be reduced to cognitive uncertainty (Bradley and Drechsler 2014). In either case, there is nothing particularly worrying about evaluative uncertainty.

However plausible this picture may be for individual decision making, the situation is less straightforward for policy making. Even on the assumption that each stakeholder has no uncertainty over their own preferences, these are not usually the same for everyone. Thus, there will be a variety of alternative preferences available to the decision maker, who may then be in a situation of evaluative uncertainty – which preferences should be used to evaluate the outcomes for the purposes of the decision: the decision maker's, some stakeholder's, or some aggregation of all preferences? And if so, which one?

Once again, the complexity of the situation brought about by the pandemic makes this uncertainty particularly severe. As we have seen, policy actions tend to impact many and various dimensions of the social world, and their consequences are not necessarily positive in all of them. This means that policy implementation will lead to important trade-offs, and stakeholders that are positioned differently with respect to these trade-offs will evaluate the outcomes of the policy differently. For the sake of illustration, we can imagine that policies protecting public health while halting the economy may be judged more favourably by the elderly, who are more exposed to the risks of the virus (Mueller, McNamara, and Sinclair 2020) while potentially enjoying more secure financial situations, than by the youth, who tend to be less exposed to COVID-19 complications than to professional insecurity.

4.5 Two arguments for stakeholders inclusion

We have seen three different types of uncertainty that play a role in decision making. All three types are made particularly severe by the circumstances of the pandemic. This severity is due not only to our ignorance of crucial information regarding the virus, but also to the complexity of our current situation – which is constituted of many interconnected dimensions – and to the presence of a plurality of stakeholders. In this section, I will argue that all three types of uncertainty provide reasons for increasing the inclusion of stakeholders in decision making. First, I suggest a democratic argument based on model and evaluative uncertainty; then, I propose an epistemic argument based on cognitive uncertainty.

A democratic argument

Under model uncertainty, the agent is uncertain about the way they have modelled their decision problem, for instance because they are not sure whether they have included all the options at their disposal or because they fear they have overlooked some relevant aspect of the situation. Under evaluative uncertainty, the agent is uncertain about the value to assign to the consequences they expect from the performance of the different actions. These states of uncertainty may not concern the agent's beliefs over some matters of fact and cannot be captured by probability distributions.

This means that cognitive uncertainty differs from model and evaluative uncertainty at least in one important respect. Insofar as something either is or is not the case, then there is a sense in which there is a correct solution to cognitive uncertainty. On the other hand, there is no unique solution to model or evaluative uncertainty, as there is no unique way to model a decision or to evaluate some consequences.

As for model uncertainty, different agents may disagree on whether something is relevant to the decision or whether some action is really available. Within some limits, this disagreement could be irreconcilable, in the sense that it may ultimately boil down to disagreement over values and priorities. The same goes for evaluative uncertainty: people may disagree on how to evaluate some outcome, even more so given the amount of trade-offs required by the complex policies responding to the pandemic. There is no system of values that is the uniquely correct one to evaluate such consequences.

While one could argue that there is no reason why, on matters of values and priorities, some viewpoints should hold more influence than others, lack of inclusion will typically favour the more powerful. In February 2020, the Italian city of Bergamo was severely hit by the pandemic. The size of the tragedy was partly due to late adoption of restrictive measures. In those crucial weeks, Bergamo's influential Confederation of Industry (Confindustria) started a campaign pressuring the city not to stop its activities, and lobbied politicians to avoid lockdowns (Horowitz and Bucciarelli 2020). In a situation of uncertainty over the value of their alternatives, decision makers under time pressure adopted the evaluative stance of the most influential stakeholder.

Confindustria's position was not wrong: it reflected a system of values. Indeed, evaluative differences may be due to different relations to the matter at stake. As each stakeholder has a different perspective on the decision and on its results, the inclusion of more stakeholders in the decision making process implies a more comprehensive and diversified perspective on the matter – one that better represents its complexity. For this reason, model and evaluative uncertainty point to the importance of including all relevant stakeholders in policy decision making.

An epistemic argument

While democratic considerations in the face of model and evaluative uncertainty call for inclusion, in the case of cognitive uncertainty epistemic considerations seem to push in the opposite direction. If there is a correct answer to our doubts, then we should not treat all possible solutions equally. Our priority should be to find the correct one, and therefore the guiding principle in assessing the alternative ways to solve the uncertainty should be one of accuracy. For this reason, addressing cognitive uncertainty is usually the role of experts and scientific advisors in policy making. However, the exclusionary nature of expertise provides an epistemic argument for the inclusion of stakeholders in decision making.

The attainment of epistemic authority requires long years of education and access to epistemic institutions. However, these two conditions are not open to all individuals equally. Education costs time and money, and people from low-income backgrounds may not have the possibility to devote resources to it. Academic institutions are by no means immune to structural inequalities (Davies and Zarifa 2012; Gillborn and Mirza 2000), and access to quality education is more difficult for minorities and under-represented groups (Carnevale and Strohl 2013). Thus, epistemic authority is not equally distributed among society, and expertise tends to be exclusionary.

However, this is not just democratically problematic – it is also epistemically problematic. The COVID-19 pandemic does not affect everybody equally (Paton, Fooks, et al. 2020; Tai, Shah, et al. 2021): it hits disproportionally ethnic (Kirby 2020; Millett, Jones, et al. 2020) and sexual (D. Banerjee and Nair 2020) minorities. These communities are more exposed not only to the disease, but also to the psychological (Pedrosa, Bitencourt, et al. 2020; Suen, Chan, and Wong 2020) and economic (Hu 2020; Kantamneni 2020) consequences of the pandemic. If this is so, then their members hold an important epistemic vantage point on the pandemic and its impacts.

By itself, diversity in science and in scientific institutions is an epistemic value (Anderson 2006; Intemann 2009; Rolin 2019), as it fosters dissent and creativity (Solomon 2006), distribution of labour (Kitcher 2003), and fights epistemic injustice (Fricker 2007). Moreover, feminist epistemology has highlighted that knowers are socially located, standing in particular relations to what is known (Anderson 2020; Harding 1986). The under-representation of minorities in science leads to the exclusion of people with first-person access to the experience of COVID-19 and of its effects: our collective enterprise of knowledge production with respect to the pandemic lacks the perspective of some of the most hit communities. For this reason, in the face of cognitive uncertainty, decision makers should strive to include a variety of stakeholders in their deliberations, and especially those whose epistemic perspective on the pandemic is particularly relevant for policy making, while also being traditionally excluded from institutional knowledge production.

One such community is represented by disabled people, who often find higher education to be an unwelcoming environment for them (Dolmage 2017; Kimball, Wells, et al. 2016). Yet, they are not only among the communities that have been most affected by the pandemic – they also hold a vantage standpoint to understand the shortcomings of the healthcare system and the chronic consequences of viral infections, due to their personal experience. And indeed, disabled people started talking about the long-term risks of mild COVID-19 infections and of the inadequacy of the institutions in front of the potential insurgence of COVID-related disabilities from the early days of the pandemic (Barbarin 2021; Marie 2022) – that is, long before the problems posed by the wave of Long Covid cases. However, their voices remained excluded from the processes of our collective decision making, when their inclusion could have provided society with the tools to better prepare to face the complexities of Long Covid.

4.6 Conclusions

The COVID-19 pandemic has shaken our lives. It presents us with a continuous series of novel problems, requiring often drastic solutions. Policy making in response to this situation may be particularly difficult, especially given the uncertainty under which we have to make decisions. However, the severity of the uncertainty cannot be a reason for inaction, as it has happened before (Oreskes and Conway 2011). Instead, looking at the nature and severity of this uncertainty may suggest some principles to guide decision making.

First, policy making under severe uncertainty should be adaptive (Walker, Rahman, and Cave 2001), to accommodate changes as the situation develops. This means that policies should have substantial room for feedback loops, to adjust their course of action as new information comes up.

Second, the presence of trade-offs in the evaluation of consequences suggests the need of hedging and mitigating measures. Mitigation aims at the reduction of the negative consequences that are sure to occur, while hedging aims at the reduction of the impact of those that might come up (Helgeson 2020). Implementing a system of unemployment benefits while imposing a lock-down is a mitigating strategy. Expanding the ICU capacity while reopening some activities is hedging against the possibility of a new wave.

Hedging and mitigating are also helpful to buy the time necessary to dissipate more cognitive uncertainty. Science needs time to provide reliable data and usable models. The more we slow down the wave of a crisis, the more we are able to attune our response to the circumstances.

Finally, I have argued that, even though the pressure to produce timely responses to the crisis may lead to swifter and less inclusive decision procedures, the inclusion of stakeholders in decision making processes is important for both democratic and epistemic reasons. This is so because agents can be uncertain not only about some factual issues, but also about the best way to model their decision problems, as well as about the value of the possible outcomes of different policy actions. Different stakeholders may have irreconcilable disagreements about these issues, which underline the importance of including all perspectives in our collective decision making. But different stakeholders may also hold different standpoints in relation to the pandemic, and their inclusion is therefore epistemically necessary to resolve our cognitive uncertainty – especially since the communities that are affected most by the pandemic are also those that are traditionally excluded by our knowledge production institutions.

Bibliography

Albert, David and Barry Loewer (1988), "Interpreting the many worlds interpretation", *Synthese*, pp. 195-213.

Alvarez, Maria (2017), "Reasons for Action: Justification, Motivation, Explanation", in *The Stanford Encyclopedia of Philosophy*, ed. by Edward N. Zalta, Winter 2017, Metaphysics Research Lab, Stanford University.

Anderson, Elizabeth (2006), "The epistemology of democracy", *Episteme*, 3, pp. 8-22.

— (2020), "Feminist Epistemology and Philosophy of Science", in *The Stanford Encyclopedia of Philosophy*, ed. by Edward N. Zalta, Spring 2020, Metaphysics Research Lab, Stanford University.

Andreoletti, Mattia and David Teira (2016), "Statistical evidence and the reliability of medical research", in *The Routledge Companion to Philosophy of Medicine*, Routledge, pp. 232-241.

Bambrough, Renford (2020), *Moral scepticism and moral knowledge*, Routledge.

Banerjee, Abhijit V., Rukmini Banerji, James Berry, Esther Duflo, Harini Kannan, Shobhini Mukerji, Marc Shotland, and Michael Walton (2017), "From proof of concept to scalable policies: Challenges and solutions, with an application", *Journal of Economic Perspectives*, 31, pp. 73-102.

Banerjee, Abhijit V., Rukmini Banerji, James Berry, Esther Duflo, Harini Kannan, Shobhini Mukherji, Marc Shotland, and Michael Walton (2016), *Mainstreaming an effective intervention: Evidence from randomized evaluations of "Teaching at the Right Level" in India*, tech. rep., National Bureau of Economic Research.

Banerjee, Abhijit V., Rukmini Banerji, Esther Duflo, Rachel Glennerster, and Stuti Khemani (2010), "Pitfalls of participatory programs: Evidence from a randomized evaluation in education in India", *American Economic Journal: Economic Policy*, 2, pp. 1-30.

Banerjee, Abhijit V., Shawn Cole, Esther Duflo, and Leigh Linden (2007), "Remedying education: Evidence from two randomized experiments in India", *The Quarterly Journal of Economics*, 122, pp. 1235-1264.

Banerjee, Abhijit V. and Esther Duflo (2011), *Poor economics: A radical rethinking of the way to fight global poverty*, Public Affairs.

Banerjee, Abhijit V., Esther Duflo, and Michael Kremer (2016), "The influence of randomized controlled trials on development economics research and on development policy", *The state of Economics, the state of the world*, pp. 482-488.

Banerjee, Debanjan and Vasundharaa S. Nair (2020), "The untold side of COVID-19: Struggle and perspectives of the sexual minorities", *Journal of Psychosexual Health*, 2, pp. 113-120.

Barbarin, Imani (2021), "Death by a thousand words: COVID-19 and the pandemic of ableist media", *Refinery29*.

Björnsson, Gunnar and Tristram McPherson (2014), "Moral attitudes for non-cognitivists: Solving the specification problem", *Mind*, 123, pp. 1-38.

Borlund, Pia (2003), "The concept of relevance in IR", *Journal of the American Society for information Science and Technology*, 54, pp. 913-925.

Bradley, Richard (2004), "Ramsey's representation theorem", *Dialectica*, 58, pp. 483-497.

— (2009), "Revising incomplete attitudes", *Synthese*, 171, pp. 235-256.

— (2017), *Decision theory with a human face*, Cambridge University Press.

Bradley, Richard and Mareile Drechsler (2014), "Types of uncertainty", *Erkenntnis*, 79, pp. 1225-1248.

Brogaard, Berit and Joe Salerno (2019), "Fitch's Paradox of Knowability", in *The Stanford Encyclopedia of Philosophy*, ed. by Edward N. Zalta, Fall 2019, Metaphysics Research Lab, Stanford University.

Broome, John (1991), *Weighing goods: Equality, uncertainty and time*, John Wiley & Sons.

— (2013), *Rationality through reasoning*, John Wiley & Sons.

Carnevale, Anthony P. and Jeff Strohl (2013), "Separate & unequal: How higher education reinforces the intergenerational reproduction of white racial privilege".

Chateauneuf, Alain and José Heleno Faro (2009), "Ambiguity through confidence functions", *Journal of Mathematical Economics*, 45, pp. 535-558.

Cohen, L. Jonathan (1994), "Some steps towards a general theory of relevance", *Synthese*, 101, pp. 171-185.

Crupi, Vincenzo (2021), "Confirmation", in *The Stanford Encyclopedia of Philosophy*, ed. by Edward N. Zalta, Spring 2021, Metaphysics Research Lab, Stanford University.

Dancy, Jonathan (2000a), *Practical Reality*, Oxford University Press, USA.

— (2000b), *Practical reality*, Oxford University Press, USA.

Davidson, Paul (1996), "Reality and economic theory", *Journal of Post Keynesian Economics*, 18, pp. 479-508.

Davies, Scott and David Zarifa (2012), "The stratification of universities: Structural inequality in Canada and the United States", *Research in Social Stratification and Mobility*, 30, pp. 143-158.

Dede, Çağlar O. (2019), "Behavioral policies and inequities: the case of incentivized smoking cessation policies", *Journal of Economic Methodology*, 26, pp. 272-289.

Dennett, Daniel C. (1987), *The Intentional Stance*, MIT Press.

Dequech, David (2000), "Fundamental uncertainty and ambiguity", *Eastern economic journal*, 26, pp. 41-60.

— (2004), "Uncertainty: individuals, institutions and technology", *cambridge Journal of Economics*, 28, pp. 365-378.

Dietrich, Franz and Brian Jabarian (2018), "Decision under normative uncertainty", *Economics & Philosophy*, pp. 1-23.

Dietrich, Franz and Christian List (2013), "A reason-based theory of rational choice", *Nous*, 47, pp. 104-134.

Dolmage, Jay T. (2017), *Academic ableism: Disability and higher education*, University of Michigan Press.

Douglas, Heather (2000), "Inductive risk and values in science", *Philosophy of science*, 67, pp. 559-579.

— (2009), *Science, policy, and the value-free ideal*, University of Pittsburgh Pre.

Dow, Sheila C. (2012), "Uncertainty about uncertainty", in *Foundations for new economic thinking*, Springer, pp. 72-82.

Edgington, Dorothy (1985), "The paradox of knowability", *Mind*, 94, pp. 557-568.

Einhorn, Hillel J. and Robin M. Hogarth (1986), "Decision making under ambiguity", *Journal of business*, S225-S250.

Ellsberg, Daniel (1961), "Risk, ambiguity, and the Savage axioms", *The quarterly journal of economics*, pp. 643-669.

Eriksson, John and Ragnar Francén Olinder (2016), "Non-cognitivism and the classification account of moral uncertainty", *Australasian Journal of Philosophy*, 94, pp. 719-735.

Fishburn, Peter (1994), "A variational model of preference under uncertainty", *Journal of Risk and Uncertainty*, 8, pp. 127-152.

Floridi, Luciano (2008), "Understanding epistemic relevance", *Erkenntnis*, 69, pp. 69-92.

— (2015), "The politics of uncertainty", *Philosophy & Technology*, 28, pp. 1-4.

Foot, Philippa (1967), "The problem of abortion and the doctrine of the double effect", *Oxford Review*, 5, pp. 5-15.

Fricker, Miranda (2007), *Epistemic injustice: Power and the ethics of knowing*, Oxford University Press.

Frigg, Roman and James Nguyen (2017), "Models and representation", in *Springer handbook of model-based science*, Springer, pp. 49-102.

Galaabaatar, Tsogbadral and Edi Karni (2013), "Subjective expected utility with incomplete preferences", *Econometrica*, 81, pp. 255-284.

Gärdenfors, Peter and Nils-Eric Sahlin (1983), "Decision making with unreliable probabilities", *British Journal of Mathematical and Statistical Psychology*, 36, pp. 240-251.

Ghirardato, Paolo (2001), "Coping with ignorance: unforeseen contingencies and non-additive uncertainty", *Economic theory*, 17, pp. 247-276.

Gilboa, Itzhak and David Schmeidler (2001), *A Theory of Case-Based Decisions*, Cambridge University Press.

Gillborn, David and Heidi Safia Mirza (2000), "Educational inequality: Mapping race, class and gender: A synthesis of research evidence".

Grant, Simon and John Quiggin (2008), *Bounded rationality and small worlds*, tech. rep., School of economics, University of Queensland.

Grzankowski, Alex (2015), "Not all attitudes are propositional", *European Journal of Philosophy*, 23, pp. 374-391.

Hansson, Sven Ove (1996), "Decision making under great uncertainty", *Philosophy of the social sciences*, 26, pp. 369-386.

Hansson, Sven Ove and Gertrude Hirsch Hadorn (2016), "Introducing the argumentative turn in policy analysis", in *The Argumentative Turn in Policy Analysis*, Springer, pp. 11-35.

Harding, Sandra G. (1986), *The science question in feminism*, Cornell University Press.

Helgeson, Casey (2020), "Structuring decisions under deep uncertainty", *Topoi*, 39, pp. 257-269.

Horowitz, Jason and Fabio Bucciarelli (2020), "The lost days that made Bergamo a coronavirus tragedy", *New York Times*.

Hu, Yang (2020), "Intersecting ethnic and native–migrant inequalities in the economic impact of the COVID-19 pandemic in the UK", *Research in Social Stratification and Mobility*, 68, p. 100528.

Huvenes, Torfinn Thomesen (2014), "Disagreement without error", *Erkenntnis*, 79, pp. 143-154.

Iacona, Andrea (2008), "Faultless or disagreement", *Relative truth*, pp. 287-295.

Intemann, Kristen (2009), "Why diversity matters: Understanding and applying the diversity component of the National Science Foundation's broader impacts criterion", *Social Epistemology*, 23, pp. 249-266.

J-PAL (2022), *About Us*, `https://www.povertyactionlab.org/about-j-pal`.

Jeffrey, Richard C. (1965), *The Logic of Decision*, University of Chicago Press.

Joyce, James M. (1999), *The foundations of causal decision theory*, Cambridge University Press.

Joyce, James M. (2010), "A defense of imprecise credences in inference and decision making", *Philosophical perspectives*, 24, pp. 281-323.

Joyce, Richard (2021), "Moral Anti-Realism", in *The Stanford Encyclopedia of Philosophy*, ed. by Edward N. Zalta, Winter 2021, Metaphysics Research Lab, Stanford University.

Kahneman, Daniel and Amos Tversky (1982), "Variants of uncertainty", *Cognition*, 11, pp. 143-157.

Kalra, Aparna (2019), "How ability-based segregation in Delhi government schools harms peer learning".

Kantamneni, Neeta (2020), "The impact of the COVID-19 pandemic on marginalized populations in the United States: A research agenda", *Journal of vocational behavior*, 119, p. 103439.

Karni, Edi and Marie-Louise Vierø (2017), "Awareness of unawareness: A theory of decision making in the face of ignorance", *Journal of Economic Theory*, 168, pp. 301-328.

Keynes, John Maynard (1921), *A treatise on probability*, Dover Publications.

Khosrowi, Donal (2019), "Trade-offs between epistemic and moral values in evidence-based policy", *Economics & Philosophy*, 35, pp. 49-78.

Khosrowi, Donal and Julian Reiss (2019), "Evidence-based policy: The tension between the epistemic and the normative", *Critical Review*, 31, pp. 179-197.

Kimball, Ezekiel W., Ryan S. Wells, Benjamin J. Ostiguy, Catherine A. Manly, and Alexandra A. Lauterbach (2016), "Students with disabilities in higher education: A review of the literature and an agenda for future research", *Higher education: Handbook of theory and research*, pp. 91-156.

Kirby, Tony (2020), "Evidence mounts on the disproportionate effect of COVID-19 on ethnic minorities", *The Lancet Respiratory Medicine*, 8, pp. 547-548.

Kitcher, Philip (2003), *Science, truth, and democracy*, Oxford University Press.

Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji (2005), "A smooth model of decision making under ambiguity", *Econometrica*, 73, pp. 1849-1892.

Knight, Frank Hyneman (1921), *Risk, uncertainty and profit*, Houghton Mifflin, vol. 31.

Korsgaard, Christine M. (1996), *The Sources of Normativity*, Cambridge University Press.

Kozyreva, Anastasia and Ralph Hertwig (2021), "The interpretation of uncertainty in ecological rationality", *Synthese*, 198, pp. 1517-1547.

Krämer, Stephan (2017), "A hyperintensional criterion of irrelevance", *Synthese*, 194, pp. 2917-2930.

Lawson, Tony (1988), "Probability and uncertainty in economic analysis", *Journal of post Keynesian economics*, 11, pp. 38-65.

Lempert, Richard O. (1976), "Modeling relevance", *Mich. L. Rev.*, 75, p. 1021.

Levi, Isaac (1983), *The enterprise of knowledge: An essay on knowledge, credal probability, and chance*, MIT press.

List, Christian (2019), *Why free will is real*, Harvard University Press.

Lombardi, Olimpia and Mario Castagnino (2008), "A modal-Hamiltonian interpretation of quantum mechanics", *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 39, pp. 380-443.

Longino, Helen E. (1990), *Science as social knowledge: Values and objectivity in scientific inquiry*, Princeton University Press.

Loomes, Graham and Robert Sugden (1986), "Disappointment and dynamic consistency in choice under uncertainty", *The Review of Economic Studies*, 53, pp. 271-282.

Luce, R. Duncan and Howard Raiffa (1957), *Games and decisions: Introduction and critical survey*, Dover Publications.

Makins, Nicholas (2021), "Attitudinal ambivalence: moral uncertainty for non-cognitivists", *Australasian Journal of Philosophy*, pp. 1-15.

Mamou, Hadrien (2020), "Rational preferences and reindividuation of relevant alternatives in decision theory: towards a theory of representation", *Topoi*, 39, pp. 283-292.

Mares, Edwin D. and Robert K. Meyer (2017), "Relevant logics", *The Blackwell guide to philosophical logic*, pp. 280-308.

Marie, Christa (2022), "Disabled people don't want to say we told you so about the pandemic, but..." *The Mighty*.

McGeer, Victoria (2007), “The Regulative Dimension of Folk Psychology”, in *Folk Psychology Re-Assessed*, ed. by Daniel D. Hutto and Matthew Ratcliffe, Kluwer/Springer Press, pp. 137-156.

Millett, Gregorio A., Austin T. Jones, David Benkeser, Stefan Baral, Laina Mercer, Chris Beyrer, Brian Honermann, Elise Lankiewicz, Leandro Mena, Jeffrey S. Crowley, et al. (2020), “Assessing differential impacts of COVID-19 on black communities”, *Annals of epidemiology*, 47, pp. 37-44.

Mueller, Amber L., Maeve S. McNamara, and David A. Sinclair (2020), “Why does COVID-19 disproportionately affect older people?”, *Aging*, 12, p. 9959.

Nielsen, Michael and Rush T. Stewart (2020), “Persistent disagreement and polarization in a Bayesian setting”, *The British Journal for the Philosophy of Science*.

O’Connor, Timothy and Christopher Franklin (2021), “Free Will”, in *The Stanford Encyclopedia of Philosophy*, ed. by Edward N. Zalta, Spring 2021, Metaphysics Research Lab, Stanford University.

Ok, Efe A., Pietro Ortoleva, and Gil Riella (2012), “Incomplete preferences under uncertainty: Indecisiveness in beliefs versus tastes”, *Econometrica*, 80, pp. 1791-1808.

Oreskes, Naomi and Erik M. Conway (2011), *Merchants of doubt: How a handful of scientists obscured the truth on issues from tobacco smoke to global warming*, Bloomsbury Publishing USA.

Pardo, Michael S. (2013), “The nature and purpose of evidence theory”, *Vand. L. Rev.*, 66, p. 547.

Paton, Alexis, Gary Fooks, Gaja Maestri, and Pam Lowe (2020), “Submission of evidence on the disproportionate impact of COVID-19, and the UK government response, on ethnic minorities and women in the UK”.

Paul, L. A. and John Quiggin (2018), “Real world problems”, *Episteme*, 15, pp. 363-382.

Pedrosa, Ana Luisa, Letícia Bitencourt, Ana Cláudia Fontoura Fróes, Maria Luíza Barreto Cazumbá, Ramon Gustavo Bernardino Campos, Stephanie Bruna Camilo Soares de Brito, and Ana Cristina Simões e Silva (2020), “Emotional, behavioral, and psychological impact of the COVID-19 pandemic”, *Frontiers in Psychology*, 11.

Perlman, Mark and Charles R. Jr McCann (1996), "Varieties of uncertainty", *Uncertainty in economic thought*, 9.

Peter, Fabienne (2021), "Truth and uncertainty in political justification", in *Political Epistemology*, ed. by Elizabeth Edenberg and Michael Hannon, Oxford University Press, pp. 64-75.

Pettigrew, Richard (2019), *Choosing for changing selves*, Oxford University Press.

Pope, Robin (1995), "Towards a more precise decision framework", *Theory and Decision*, 39, pp. 241-265.

Pratham (2019), *Pratham and J-PAL awarded Co-Impact grant for Africa initiative*, `https://prathamusa.org/press/pratham-and-j-pal-awarded-co-impact-grant-for-africa-initiative/`.

Pravato, Giulia (2020), "Indeterminacy and Normativity", *Erkenntnis*, pp. 1-23.

Ramsey, F. (1926), "Truth and probability", in *Philosophy of Probability: Contemporary Readings*, ed. by Antony Eagle, Routledge, pp. 52-94.

Ravallion, Martin (2014), "Taking ethical validity seriously", *Development Impact*.

— (2020), *Should the randomistas (continue to) rule?*, tech. rep., National Bureau of Economic Research.

Raz, Joseph (1999), *Practical reason and norms*, Oxford University Press.

Reed, Baron (2021), "Certainty", in *The Stanford Encyclopedia of Philosophy*, ed. by Edward N. Zalta, Fall 2021, Metaphysics Research Lab, Stanford University.

Rolin, Kristina (2019), "The epistemic significance of diversity", *The Routledge Handbook of Social Epistemology*.

Rudner, Richard (1953), "The scientist qua scientist makes value judgments", *Philosophy of science*, 20, pp. 1-6.

Rumsfeld, Donald H. (2002), "Department of Defense news briefing".

Saunders, Simon (2010), "Chance in the Everett interpretation", in *Many worlds? Everett, quantum theory, & reality*, ed. by Simon Saunders, Jonathan Barrett, Adrian Kent, and David Wallace, Oxford University Press, pp. 181-205.

Savage, Leonard J. (1954), *The foundations of statistics*, Courier Corporation.

Scanlon, Thomas (1998), *What we owe to each other*, Harvard University Press.

Schipper, Burkhard C. (2014), "Unawareness— A gentle introduction to both the literature and the special issue", *Mathematical Social Sciences*, 70, pp. 1-9.

Schroeder, Mark (2007), *Slaves of the Passions*, Oxford University Press.

Shafer, Glenn (1986), "Savage revisited", *Statistical science*, pp. 463-485.

Shafer-Landau, Russ (1994), "Ethical disagreement, ethical objectivism and moral indeterminacy", *Philosophy and Phenomenological Research*, 54, pp. 331-344.

Sher, Itai (2019), "Comparative value and the weight of reasons", *Economics & Philosophy*, 35, pp. 103-158.

Singh, Pritam Pal (2019), "Delhi: Chunauti scheme faces challenge in Delhi High Court".

Solomon, Miriam (2006), "Norms of epistemic diversity", *Episteme*, 3, pp. 23-36.

Sorensen, Roy (2018), "Vagueness", in *The Stanford Encyclopedia of Philosophy*, ed. by Edward N. Zalta, Summer 2018, Metaphysics Research Lab, Stanford University.

Sperber, Dan and Deirdre Wilson (1986), *Relevance: Communication and cognition*, Citeseer, vol. 142.

Steel, Daniel (2010), "Epistemic values and the argument from inductive risk", *Philosophy of science*, 77, pp. 14-34.

Steele, Katie and H. Orri Stefánsson (2020), "Decision Theory", in *The Stanford Encyclopedia of Philosophy*, ed. by Edward N. Zalta, Winter 2020, Metaphysics Research Lab, Stanford University.

— (2021), *Beyond Uncertainty: Reasoning with Unknown Possibilities*, Cambridge University Press.

Suen, Yiu Tung, Randolph C. H. Chan, and Eliz Miu Yin Wong (2020), "Effects of general and sexual minority-specific COVID-19-related stressors on the mental health of lesbian, gay, and bisexual people in Hong Kong", *Psychiatry research*, 292, p. 113365.

Tai, Don Bambino Geno, Aditya Shah, Chyke A. Doubeni, Irene G. Sia, and Mark L. Wieland (2021), "The disproportionate impact of COVID-19 on racial and ethnic minorities in the United States", *Clinical Infectious Diseases*, 72, pp. 703-706.

Tersman, Folke (2006), *Moral disagreement*, Cambridge University Press.

Thoma, Johanna (2019), "Decision Theory", in *The Open Handbook of Formal Epistemology*, ed. by Richard Pettigrew and Jonathan Weisberg, PhilPapers Foundation, pp. 57-106.

Tversky, Amos and Daniel Kahneman (1989), "Rational choice and the framing of decisions", in *Multiple criteria decision making and risk analysis using microcomputers*, Springer, pp. 81-126.

Van Benthem, Johan (2004), "What one may come to know", *Analysis*, 64, pp. 95-105.

Van Ditmarsch, Hans, Wiebe Van der Hoek, and Petar Iliev (2012), "Everything is knowable–how to get to know whether a proposition is true", *Theoria*, 78, pp. 93-114.

van Roojen, Mark (2018), "Moral Cognitivism vs. Non-Cognitivism", in *The Stanford Encyclopedia of Philosophy*, ed. by Edward N. Zalta, Fall 2018, Metaphysics Research Lab, Stanford University.

von Neumann, John and Oskar Morgenstern (1944), *Theory of Games and Economic Behavior*, Princeton University Press.

Walker, Warren E., S. Adnan Rahman, and Jonathan Cave (2001), "Adaptive policies, policy analysis, and policy-making", *European journal of operational Research*, 128, pp. 282-289.

Weirich, Paul (1986), "Expected utility and risk", *The British journal for the philosophy of science*, 37, pp. 419-442.

Williams, Bernard (1979), "Internal and external reasons", in *Rational Action*, ed. by Ross Harrison, Cambridge University Press, pp. 101-113.

Winkler, Robert L. (1996), "Uncertainty in probabilistic risk assessment", *Reliability Engineering & System Safety*, 54, pp. 127-132.

Zuckerman, Diana M. (2006), "FDA Advisory Committees: Does approval mean safety?"