This is the authors' final version; the article has been accepted for publication in *Perspectives on Science*

Non-empirical uncertainties in evidence-based decision making.

Malvina Ongaro (University of Eastern Piedmont)

Malvina.ongaro@uniupo.it

Mattia Andreoletti (Vita-Salute San Raffaele University)

Andreoletti.mattia@unisr.it

Abstract. The increasing success of the evidence-based policy movement is raising the demand of 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 start by introducing a distinction between empirical and non-empirical types of uncertainty, and we explore the role of two non-empirical uncertainties in the context of policy making. We argue that the authority of scientific advisors is limited to empirical uncertainty and cannot extend beyond it. While the appeal of evidence-based policy rests on a view of scientific advice as limited to empirical 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-empirical types of uncertainty.

This is the authors' final version; the article has been accepted for publication in *Perspectives on Science*

1. Evidence-based decision making.

Uncertainty is a pervasive feature of decision making. Decisions are aimed at reaching or avoiding some results, 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, Banerjee & Duflo (2011: 16) 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 (Banerjee et al. 2016a; Ravallion 2018), culminating in the recent Nobel Prize awarded to some of its most prominent figures (Royal Swedish Academic of Sciences 2019).

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 focussed on the role of

values in evidence gathering (e.g. Rudner 1953; Longino 1990; Douglas 2000; Steel 2010), the rise of EBP has fostered new sets of questions. Some authors have questioned the ethical validity of randomisation (Ravallion 2014; 2018), 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 & 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 listen 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-empirical uncertainty, building on Bradley & Drechsler's (2014) seminal work on the taxonomy of uncertainties. We focus on two types of non-empirical 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 empirical 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.

2. Non-empirical 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 2020), 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 (e.g. Ramsey 1926; 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. Levi 1980; Joyce 2010) or second-order

weights on probabilities (e.g. Gärdenfors & Sahlin 1983; Klibanoff et al. 2005; Chateauneuf & Faro 2009).

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 was provided. In principle, there is some evidence that could effectively resolve the agent's practical uncertainty. 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. her beliefs about what is or will be the case. Following Bradley & Drechsler (2014), we call this type of uncertainty *empirical.* However, there is growing work showing that not all the uncertainty is empirical. For instance, in his typology of uncertainty, Hansson (1996) identifies four components of what he labels "great" uncertainty, neither of which concerns the state of the world. Helgeson (2020) investigates "deep" uncertainty in the structuring of decisions. Bradley & Drechsler (2014) themselves identify uncertainties concerning normative and modal judgements, as well as descriptive (see also Dietrich & Jabarian 2018). *Non-empirical* uncertainties are those types of uncertainty that do not concern the actual state of the world, and that therefore cannot be solved by evidence.

If there are non-empirical 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 accommodated, one way or another. If, as we say, some of this uncertainty is insensitive to evidence because it is non-empirical, then no policy can ever be chosen entirely on evidential grounds. Choosing a policy will always require some non-empirical judgements, no matter the amount of evidence available.

In the next two sections, we focus on two types of non-empirical uncertainty that play an important role in decisions for policy making. We build on Bradley & Drechsler's taxonomy and identify these uncertainties as *ethical* and *state space* uncertainty. In each case, we present the uncertainty starting from Bradley & Drechsler's account. Then, we clarify the sense in which these uncertainties are non-empirical. Finally, we elaborate on the original account to illuminate the elements relevant to policy making.

2.1 Ethical uncertainty

The agent picks a certain option when she finds it more desirable than the alternatives. If she is uncertain with respect to the desirability of some alternative, she is in a situation of *ethical uncertainty*. According to Bradley & Drechsler, ethical uncertainty arises when the values used to assess the desirability of the different alternatives are unknown or non-existent (2014: 1237). They identify three main views on the nature and possibility of this uncertainty, and we will show that each has different implications for policy making.

The first position is ethical subjectivism, which is taken to be the most prominent position. According to this view, ethical uncertainty concerns the degree to which a certain alternative is desirable, and so it concerns one's own judgements of desirability. For the

subjectivist, ethical uncertainty is virtually impossible, as it would require one to be uncertain about their own tastes. The opposite view is labelled ethical cognitivism. This view takes ethical uncertainty to be about beliefs over some objective normative facts. It is a matter of truth whether something is good or not, and therefore desirability can be the object of uncertain beliefs. The third view they consider to be intermediate between the first two. It takes ethical uncertainty to concern judgements over some factual properties of the agent, namely her tastes.

Now, the last two views reduce ethical uncertainty to a case of empirical uncertainty. For the cognitivist, ethical uncertainty is simply empirical uncertainty over a specific class of beliefs. However, the reduction of ethical uncertainty to empirical uncertainty comes at the cost of assuming realism about what is good. In the intermediate view, ethical uncertainty concerns factual judgements, and therefore is empirical - but it requires to see tastes as facts about someone.

On the other hand, the ethical subjectivist does not see tastes as some fact about what is desirable or about the agent. If she 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.

Without trying to settle the debate over which of these views is the adequate understanding of ethical 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.

Ethical cognitivism implies that there is a correct way to evaluate the different alternatives. If the decision maker is uncertain, then she can find out something more about what is valuable. In sum, she should treat her uncertainty over values just as her uncertainty over facts: as something that can go away with the right information. This is so whether the decision concerns individual choices or policy options. While this position makes ethical uncertainty a fully empirical issue, it does so at the cost of requiring strong realism over what is good - a position that may be at odds with the practices of democracy, which typically assume some pluralism.

The intermediate view 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 ethical uncertainty in policy contexts amounts to uncertainty over the stakeholders' evaluation of the alternatives. According to the intermediate view, this is empirical uncertainty over some factual features of the people involved: each stakeholder has some tastes, and ethical uncertainty means not knowing some of these tastes. However, while ethical uncertainty could be reduced to empirical 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, she would still need to aggregate all that information into one single evaluation. But there is no independently correct procedure to do so. Non-empirical uncertainty comes back at this higher-level decision making.

Let us now look at the fully non-empirical 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 non-empirical.

2.2. State space uncertainty

The *state space* of a decision is the set of possible states of the world, i.e. the worldly contingencies on which the outcomes depend. Standardly, decision theory takes the state space as given, and does not question the elements included in or excluded from it. Nonetheless, in real-world decisions state spaces are not ready-made: the agent has to build her own. State space uncertainty arises whenever the decision maker is aware of the possibility that she has not included all the relevant contingencies in her state space (Bradley & Drechsler 2014: 1245).

The possibility of state space uncertainty has given rise to a prolific literature on decision making under unawareness, i.e. in those circumstances where the agent is aware that she may be unaware of some contingency (Schipper 2014; Karni & Vierø 2017). In this basic

form, this type of uncertainty seems to be empirical, albeit of extreme severity. Far from being able to assign precise probabilities, the agent has no information whatsoever regarding those contingencies of which she is unaware. However, it is a matter of fact whether those contingencies obtain or not, and in principle there is some evidence that could resolve her uncertainty.

There is, however, an important non-empirical element to state space uncertainty. It is clear that the agent cannot be aware of *all* possible contingencies, let alone include them in her state space. Indeed, state space uncertainty concerns all the *relevant* contingencies. The question that the agent faces when constructing the state space is not whether she has included all the possible things that may be the case or that may happen - rather, she faces the question of *what to include* in the state space. This means that unawareness is only one way in which some relevant contingency may be left out of the state space. The other is selection - the agent may overlook something, consciously excluding it as something that is not relevant to the matter at hand.

Since it is impossible to include every eventuality in the state space, then it is always possible that the agent has omitted something relevant. Thus, state space 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, to an important extent being relevant to a decision is not an empirical property. Relevance is the result of a subjective judgement of the decision maker. Therefore, state space uncertainty is non-empirical, because what counts as relevant is not an external fact.

If this is so - if judgements of relevance are subjective - then the question of what to include in the state space 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 consequences in the structuring of the decision problem. The agent's decision does not depend only on factual contingencies, but also on which consequences she expects from the performance of each alternative option. Now, actions can have a variety of effects. The agent may be unaware of some of these - she may not expect a certain option to result in a certain effect, and therefore exclude it from her considerations. But she may also exclude some effects as irrelevant for her 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. But their disagreement is not due to lack of

sufficient information, as to different priority settings. And priority setting has more to do with normative considerations than with empirical ones.

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 she lacks some crucial information about the actual state of the world (empirical uncertainty), because she does not know how to evaluate the possible consequences of her actions (ethical uncertainty), or because she is not sure about which contingencies and consequences she should take into consideration (state space uncertainty). In the context of policy making, the presence of a variety of stakeholders implies that both ethical uncertainty and state space 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, she will have to solve the practical uncertainty surrounding the decisions: to say that option *a* is better than option *b* she has to solve the decision problem. In order to do that, the scientific advisor has to face both empirical and non-empirical uncertainties. On the other

hand, scientific advice along the second route does not require a stance on what to do, and thus the resolution of non-empirical uncertainties.

As any stakeholder, the scientific advisor can form an opinion on the decision at stake. In order to do so, she will rely on her expertise to resolve the empirical components of the uncertainty. But as for the non-empirical components, she is 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 grant practical authority, i.e. authority over what should be done. Thus, the opinion of the scientific advisor formed on her 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 empirical 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 empirical 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 empirical 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 & Teira 2019). 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 criticizes 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 empirical 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 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 empirical uncertainty, where however 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 empirical 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. (Banerjee et al. 2017: 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 empirical uncertainty.

4. A real world application: the case of TaRL Africa

So far, we have analysed two types of non-empirical 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 2020), 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)[1], 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 (Banerjee 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 (Banerjee et al. (2007), (2010), (2016b)), TaRL reached a national scale, involving millions of kids.

In Delhi, government schools applied Pratham's methods in 2016. But TaRL has been seen as implementing a policy of segregation in classrooms, leading some parents to file a lawsuit against Delhi government[2]. In fact, teachers testify that level-based division has an impact on the identity formation of children, that are now often identified with their level. According to them, this has translated into bullyism and demotivation. Moreover, teachers accuse that, within this discriminatory system, students classified as "bad learners" are given less interesting and engaging teaching activities[3].

Banerjee 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

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

[2] Delhi: Chunauti scheme faces challenge in Delhi High Court

[3] https://thewire.in/education/delhi-schools-ability-based-grouping

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[4]. 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". (J-PAL news release).

However, the experiments provide reasons for a policy to the extent that policy choice requires the reduction of empirical 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

[4] https://www.co-impact.org/co-impact-announces-80-million-in-grants-aimed-at-improving-the-lives-of-9-million-people-in-africa-south-asia-and-latin-america

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-empirical uncertainties are at play in the TaRL case.

As we have seen, ethical 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 would 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 state space 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 subjective, then disagreements are both possible and legitimate, and cannot be settled by resorting to evidence.

5. Conclusion

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. 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 empirical and non-empirical 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 empirical uncertainty, in practice its impact on decision making tends to be larger. The case of TaRL Africa shows how EBP researchers may move beyond empirical uncertainty to promote a policy that their evidence proves to be effective. As promoting a policy implies having chosen that policy, it faces ethical and state space uncertainty. Their resolution requires judgements of relevance and desirability that are outside of the epistemic authority of scientific advisors.

Acknowledgements

We are grateful to two anonymous reviewers for their constructive comments.

This is the authors' final version; the article has been accepted for publication in *Perspectives on Science*

Mattia Andreoletti has been funded by the PRIN Project "From Models to Decisions".

References

Andreoletti, M., & Teira, D. 2016. Statistical Evidence and The Reliability of Medical Research. *The Routledge Companion to Philosophy of Medicine*, 218.

Banerjee, A., Cole, S., Duflo, E., & Linden L. 2007. Remedying Education: Evidence from Two Randomized Experiments in India. *Quarterly Journal of Economics, 122*(3): 1235–64.

Banerjee, A., Banerji, R., Duflo, E., Glennerster, R., & Khemani, S. 2010. Pitfalls of Participatory Programs: Evidence from a Randomized Evaluation in India. *American Economic Journal: Economic Policy 2*(1): 1–30.

Banerjee, A., & Duflo, E. 2011. *Poor economics: rethinking poverty & the ways to end it.* Random House.

Banerjee, A. V., Duflo, E., & Kremer, M. 2016a. The influence of randomized controlled trials on development economics research and on development policy. *The State of Economics, The State of the World*.

Banerjee, A., Banerji, R., Berry, J., Duflo, E., Kannan, H., Mukherji, S., Shotland, M., & Walton, M. 2016b. Mainstreaming an effective intervention: Evidence from randomized evaluations of "Teaching at the Right Level" in India (No. w22746). *National Bureau of Economic Research.*

Banerjee, A., Banerji, R., Berry, J., Duflo, E., Kannan, H., Mukerji, S., Shotland, M., & Walton, M. 2017. From proof of concept to scalable policies: Challenges and solutions, with an application. *Journal of Economic Perspectives, 31*(4), 73-102.

Bradley, R., & Drechsler, M. 2014. Types of uncertainty. *Erkenntnis*, *79*(6), 1225-1248.

Dede, O. Ç. 2019. Behavioral policies and inequities: the case of incentivized smoking cessation policies. *Journal of Economic Methodology*, *26*(3), 272-289.

Dietrich, F., & Jabarian, B. 2018. Decision under normative uncertainty. ⟨halshs-01903642⟩.

Chateauneuf, A., & Faro, J. H. 2009. Ambiguity through confidence functions. *Journal of Mathematical Economics*, *45*(9-10), 535-558.

Douglas, H. 2000. Inductive risk and values in science. *Philosophy of science*, *67*(4), 559-579.

Douglas, H. 2009. *Science, policy, and the value-free ideal*. University of Pittsburgh Press.

Gärdenfors, P., & Sahlin, N. E. 1983. Decision making with unreliable probabilities. *British Journal of Mathematical and Statistical Psychology*, *36*(2), 240-251.

Hansson, S. O. 1996. Decision making under great uncertainty. *Philosophy of the social sciences*, *26*(3), 369-386.

Helgeson, C. 2020. Structuring decisions under deep uncertainty. *Topoi*, *39*(2), 257-269.

J-PAL 2020. *About us.* https://www.povertyactionlab.org/about-j-pal (retrieved 12/02/2020).

Joyce, J. M. 2010. A defense of imprecise credences in inference and decision making. *Philosophical perspectives*, *24*, 281-323.

Karni, E., & Vierø, M. L. 2017. Awareness of unawareness: A theory of decision making in the face of ignorance. *Journal of Economic Theory*, *168*, 301-328.

Khosrowi, D. 2019. Trade-offs between epistemic and moral values in evidence-based policy. *Economics & Philosophy*, *35*(1), 49-78.

Khosrowi, D., & Reiss, J. 2019. Evidence-Based Policy: The Tension Between the Epistemic and the Normative. *Critical Review*, *31*(2), 179-197.

Klibanoff, P., Marinacci, M., & Mukerji, S. 2005. A smooth model of decision making under ambiguity. *Econometrica*, *73*(6), 1849-1892.

Levi, I. 1980. *The enterprise of knowledge: An essay on knowledge, credal probability, and chance*. MIT press.

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

Peter, F. 2020. Truth and Uncertainty in Political Justification. In Edenberg, E. and Hannon, M. (eds.) *Politics and Truth: New Perspectives in Political Epistemology.* Oxford: Oxford University Press (forthcoming).

Ramsey, F. P. 1926. Truth and probability. In *Readings in formal epistemology* (2016: 21-45). Springer, Cham.

Ravallion, M. 2014. Taking ethical validity seriously. *Development Impact*.

Ravallion, M. 2018. Should the randomistas (continue to) rule?. *Center for Global Development Working Paper*, 492.

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

Savage, L. J. 1954. *The foundations of statistics*. Dover: New York.

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

Steel, D. 2010. Epistemic values and the argument from inductive risk. *Philosophy of Science*, *77*(1), 14-34.

Zuckerman, D. 2006. FDA Advisory Committees: Does Approval Mean Safety?. *Washington, DC: National Research Center for Women & Families* [Online]. Available: http://www.center4research.org/ news/fda-app-safety.html [2006 August 29]