1 Introduction

The coronavirus pandemic that swept the globe in 2020 highlighted the myriad ways that values intersect with scientific research. For example, the speed with which the biomedical research community pivoted towards studying the virus and developing vaccines illustrates how important it is to be able to steer research so that it addresses what society values. Values were also obvious in policymakers' decisions about which public policies to enact in response to the pandemic. When deciding what kinds of lockdowns or mask mandates to impose, they not only had to assess the efficacy of various measures for suppressing the virus's spread but also had to weigh the overall social costs and benefits of taking those measures. At first glance, it might seem the policymakers could just 'follow the science', but it is clear upon reflection that responsible leaders had to consider a range of factors (e.g., economics, community well-being, individual rights, and mental health) when deciding how to craft their responses (Hilgartner et al. 2021). Leaders also had to take values into account when deciding how to describe the disease. For example, US President Donald Trump was criticised for referring to the 'Chinese virus' because of its tendency to promote racial stereotypes and stigmatisation (Viala-Gaudefroy and Lindaman 2020). Similarly, the World Health Organisation (WHO) faced decisions about how to describe variants of the disease, given the potential to stigmatise particular countries by naming variants after them. The WHO ultimately decided it was more responsible to name variants using Greek letters (Konings et al. 2021).

The process of developing, testing, and distributing therapies and vaccines was also awash in values. For example, the quick development of COVID-19 vaccines was partly the result of public–private partnerships, such as Operation Warp Speed in the United States, which helped fund vaccine development by private companies. However, because public funds played such a significant role in the development of these vaccines (and because the pandemic represented such a significant public-health emergency), some policymakers argued that the patents for these vaccines should be waived so that lower-income countries could more readily afford them (Iacobucci 2021).

Another important ethical debate related to vaccine development was whether to engage in human challenge trials of COVID-19 vaccines (Cornwall 2020). In this kind of trial, people are administered a vaccine or a placebo and then deliberately infected with the disease. This approach can speed the development of vaccines because the developers do not have to wait for people to be exposed to the disease by chance, but it also raises ethical concerns about the appropriateness of deliberately exposing people to a potentially deadly disease for which there are only limited treatments. Even in standard trials that did not involve deliberately exposing people to the disease, the designers had to make value-laden decisions about what populations to include. Initial trials did not include children because of a desire to protect them from harm, but eventually, companies began to design trials for children as well so the vaccines could be made available to them (Jenco 2020). Some researchers also worried that racial and ethnic minority groups were under-represented in trials of COVID-19 therapies (Chastain et al. 2020). Given that these groups were among those most severely affected by the pandemic, it was especially important that they be adequately represented in studies designed to test the safety and efficacy of treatments.

The disproportionate impacts of the pandemic on minoritised racial or ethnic groups also illustrate a plethora of broader ways in which the values of society affect the practices of science and determine how well those practices serve (or fail to serve) the interests of particular communities. For example, in countries like the United States, unequal access to economic resources and medical care contributed to a slower roll-out of COVID-19 vaccines in Black communities (Johnson et al. 2021). This was especially harmful because those same social disadvantages meant that Black people were more likely to experience pre-existing conditions that aggravated the effects of the disease (Valles 2020). Another factor depressing vaccination rates among minoritised communities was their distrust of the medical establishment, which was fuelled by past scandals and the legacies of racist science.[1] For example, in the infamous Tuskegee syphilis study that spanned four decades of the twentieth century, researchers observed the course of untreated syphilis in Black men without giving them information about available treatments (Reverby 2000). More broadly, the biological, medical, and social sciences amplified the racist and sexist values of society for centuries by trying to identify, measure, and explain differences between different races and sexes (Kendi 2016; Kourany 2010, 2020). In sum, the values of society have helped foster a context in which Black people in the United States are less likely to benefit from medical innovations, and the legacy of racist scientific and medical practices has sometimes discouraged Black people from taking advantage of those innovations even when they are available.

This brief reflection on the COVID-19 pandemic illustrates that values pervade scientific practice, that the influences of values can be obvious in some cases and obscure in others, and that those influences can be both good

[1] See, e.g.,www.pewresearch.org/fact-tank/2020/06/04/black-americans-face-higher-covid-19-risks-are-more-hesitant-to-trust-medical-scientists-get-vaccinated/

and bad. Values can steer the direction of research, influence the design of studies, affect how results are described and interpreted, and guide the ways that scientific information is used. Sometimes, the influences of values reflect conscious goals or concerns, as when ethicists debated whether to pursue human challenge trials, but they can also be much more subtle, as when the racist history of biomedical science contributed to vaccine hesitancy. Values can be beneficial, as when concerns about equity prompted researchers to scrutinise whether minoritised groups were included in trials, and they can also be harmful, as when the language for describing the pandemic contributed to the stigmatising of racial groups or geographic locations.

This Element is designed to foster a greater understanding of this complex landscape and to provide recommendations for responsibly managing the roles that values play in science. Section 2 examines how values influence science. In the process, it clarifies what we mean by 'values' and the ways in which scientific judgements can be 'value-laden'. Section 3 explores whether it is appropriate to deliberately bring values into the practice of science and, if so, under what conditions. Section 4 considers a range of proposals for responsibly managing the roles that values play in science. Finally, Section 5 proposes a path forward for those who want to help promote responsible roles for values in science. Whether values are brought into science intentionally or unintentionally, they clearly have significant influences on scientific practice. Therefore, it is important to develop thoughtful strategies for harnessing those influences so that the power of science can be directed towards the greatest social good.

2 How Do Values Influence Science?

In order to analyse the roles that values play in science, this section begins by providing some clarity about what it means to talk about 'values' and their influences on 'judgements' in science. It then provides a systematic description of the many different ways that values intersect with judgements in science. This analysis demonstrates the pervasive entanglements between science and values and illustrates how important it is to think more carefully about how to manage them.

2.1 Values

In a very basic sense, a value can be defined as something that is desirable or worthy of pursuit (Elliott 2017, 11).[2] Once one attempts to elaborate on this

[2] This Element provides a broad definition of values, but it is worth keeping in mind that some scholars have argued for providing a narrower definition of values and a clarification of their relationships to the range of other contextual factors that can influence science (see, e.g., Biddle 2013). This is an important topic for further investigation.

definition, however, a number of difficult questions arise. Are these values desirable in an objective sense that is independent of what people actually happen to desire, or are they just subjectively desirable insofar as someone actually desires them (Brown 2018)? Building on this question, do values have 'cognitive status', in the sense that one can provide evidence for their desirability (Brown 2020)? And are these values actually qualities or states of affairs out 'in the world', or do they refer to the beliefs or concepts in our minds that represent these valuable things (Schwartz and Bilsky 1987)? Finally, what are the best ways of classifying these desirable qualities (see, e.g., Brown 2020; Schwartz and Bilsky 1987; Scriven 1974)? This question is important because those writing about values in science often lament that the word 'value' is used as a label for a very wide array of phenomena that ought to be treated in different ways (Biddle 2013; Brown 2020; Rooney 2017; Ward 2021).

For the purposes of this Element, I will set aside most of these theoretical questions. However, a crucial issue that pervades the literature on science and values is whether values can be distinguished into the categories of those that are 'epistemic' (i.e., indicative of truth or knowledge) and those that are 'non-epistemic'. Assuming that science is directed at the epistemic tasks of identifying true or reliable information about the world, one might think that only the 'epistemic' values that advance these tasks have a proper role to play in scientific reasoning. For example, building on the earlier work of Thomas Kuhn (1977), Ernan McMullin (1983) argued that the qualities of predictive accuracy, internal coherence, external consistency, unifying power, and fertility should be regarded as epistemic values because they are indicators that scientific theories are true. McMullin classified other values as non-epistemic because he did not think they promoted the epistemic goals of science. Importantly, though, he admitted that it is not always obvious whether particular qualities count as epistemic values or not; for example, he argued that scientists could learn over time whether a quality like simplicity actually serves as a reason for thinking that a theory is likely to be true (McMullin 1983).

It turns out that science is awash in values, and many of these values are not clearly epistemic. In addition to the values that serve as indicators of a theory's desirability, there are values that serve as goals for individual scientists, such as advancing in one's career, gaining recognition, generating discoveries, obtaining financial resources, and serving society. There are also ethical and social values that can apply to many different aspects of science; these include goals like economic development, public health, animal welfare, environmental conservation, and social justice. There are values that apply to scientific collaborations, such as being honest, giving appropriate credit, treating others with respect, providing open access to resources, and promoting equity and inclusion

Cambridge University Press & Assessment
978-1-009-05563-5 — Values in Science
Kevin C. Elliott
Excerpt
More Information

(Rolin 2015). There are also values that apply to specific domains or kinds of science; for example, values like standardisation, reproducibility, usability, and efficiency are particularly relevant to policy-relevant fields of science.

Given the diversity and complexity of all these different kinds of values, the distinction between epistemic and non-epistemic values has proven to be highly controversial. McMullin (1983) himself acknowledged that a value could be epistemic in some contexts but not in others. Daniel Steel elaborated on this point by distinguishing ‘intrinsic’ epistemic values, which are constitutive of or necessary for truth, from ‘extrinsic’ epistemic values, which ‘promote the attainment of truth without themselves being indicators or requirements of truth’ (Steel 2010, 18). For example, Steel regards predictive accuracy and internal consistency as intrinsic epistemic values, whereas he characterises testability and simplicity as extrinsic epistemic values. Intrinsic epistemic values are always indicators of truth, but an extrinsic epistemic value like external consistency (i.e., consistency between a theory or hypothesis and a scientist’s other beliefs) could be an indicator of truth in some contexts but not in other contexts (Steel 2010). Heather Douglas (2013) has further categorised epistemic values into those that count as minimal criteria as opposed to those that are ideal desiderata, as well as those that apply directly to theories as opposed to those that apply to the relationship between theories and evidence. Minimal criteria like internal consistency and empirical adequacy must be met in order for a theory to be epistemically acceptable; in fact, some philosophers argue that it would be better to refer to these qualities as epistemic criteria rather than values (Douglas 2009; Norton 2021). In contrast, Douglas argues that some desiderata that apply to theories (e.g., having a broad scope) are not themselves indicators of truth, but they can help evaluate the truth of theories. Because of the potential for values to be pragmatically helpful for arriving at truths but not indicative of truth themselves, Douglas prefers to talk about ‘cognitive values’ rather than ‘epistemic values’ (see, e.g., Douglas 2013).

All these distinctions might appear to be somewhat pedantic, but they are important because some philosophers try to manage values in science by allowing epistemic values to influence scientific reasoning while limiting the role of non-epistemic values (see, e.g., Lacey 2017; McMullin 1983). If one cannot maintain a compelling distinction between these two kinds of values, then this avenue for managing values breaks down. For example, Helen Longino (1996) has challenged the distinction between epistemic and non-epistemic values by arguing that the typical list of epistemic values developed by Kuhn and McMullin is not as uncontroversial or value-neutral as they thought; according to Longino, it is unrealistic to think that one can identify ‘pure’ epistemic values. Along similar lines, Phyllis Rooney (1992) challenged the distinction between epistemic and non-epistemic values by showing that non-epistemic factors can shape the ways epistemic values are applied and interpreted. More recently, she has argued that many values fall into a grey area that is neither purely epistemic nor non-epistemic (Rooney 2017). Given this complexity, the categorisation of values is clearly a topic that merits further scrutiny by philosophers of science.

2.2 Value Judgements

To analyse the role of values in science, it is also crucial to clarify the concept of ‘value judgements’. This section provides an overview of what they are, what terms will be used to describe them in this Element, and what exactly it means to say that judgements are ‘value-laden’. Speaking broadly, value judgements are decisions that involve the weighing of values, but these judgements can take multiple forms (Scriven 1974). One kind of value judgement involves assessing whether and to what extent a particular quality really is desirable in a particular context. For example, how desirable is it for a theory to be simple? And how important is it for scientific data to be made publicly available? Another kind of value judgement involves assessing the extent to which a particular value has been achieved. For example, to what extent does a particular method of housing experimental animals promote their welfare? To what extent does a particular theory exhibit explanatory power? Another kind of value judgement involves weighing the importance of different values against each other. For example, is it more important for a model to be highly predictively accurate or for it to have broad applicability? Or, when one is in doubt about the severity of an environmental threat, is it better to overestimate the threat (thereby prioritising environmental conservation) or to underestimate it (thereby prioritising short-term economic development)?

Kuhn (1977) famously argued that the assessment of scientific theories involved two kinds of value judgements. First, theory assessment requires determining the extent to which particular theories exemplify particular values, like fertility or scope. Second, theory assessment involves deciding how much weight to place on different values when rival theories display them to differing extents. Thus, Kuhn emphasised that scientific reasoning is not an algorithmic, rule-governed endeavour; instead, it involves complex choices (i.e., value judgements) on which reasonable scientists can disagree.

This Element will refer to ‘judgements’ or ‘choices’ rather than ‘value judgements’, but it will often refer to these judgements as being ‘value-laden’. One reason for this terminological decision is that the language of 'value judgements' can be confusing for those who are not familiar with the philosophical literature on this topic. Those outside the philosophical community are likely to assume that 'value judgements' necessarily involve assessments of ethical or social values, whereas philosophers of science use the term more broadly to refer to any choices that involve weighing multiple desiderata in ways that are not rule-governed. Speaking of 'judgements' or 'choices' in general is less likely to cause this sort of confusion.

Referring to these judgements as 'value-laden' is a way to highlight the fact that these choices can incorporate values in subtle ways that are not immediately obvious. Zina Ward (2021) has provided a very helpful discussion of four ways in which judgements in science can be value-laden. First, values can provide *motivating reasons* for judgements, in the sense that scientists can either consciously or unconsciously decide to make judgements in particular ways because of values. Second, judgements can be value-laden in the sense that values provide *justifying reasons* in favour of making them in a particular way. Ward notes that it is important to distinguish between motivating and justifying reasons because someone could be psychologically motivated to make a decision because of one set of reasons even though there are different reasons that actually justify the decision. Third, values can serve as *causes* for a judgement even when they do not act as motivating or justifying reasons. For example, this could happen if values contribute to setting up institutions like universities or funding agencies in particular ways, and those institutional structures influence scientific judgements in indirect ways not envisioned or intended by those creating the institutions. Fourth, values can *be impacted or affected by* judgements in science, such as when the choice to use a particular study design tends to overestimate an environmental health threat, thereby prioritising public health over the short-term economic interests of those generating the threat.[3] This Element will use the term 'value-laden judgements' to describe all four of these scenarios in which values intersect with scientific choices.

2.3 Relationships between Values and Science

Building on the concepts discussed in Sections 2.1 and 2.2, it is clear that there are a wide array of judgements in science that can be value-laden. Figure 1

[3] Given the ambiguous way in which values can be thought of either as ideas in our minds or as things in the world, it is important to clarify that Ward is conceptualising values as things in the world when she says that judgements can impact or affect values. In other words, she is pointing out that judgements in science can promote some things in the world that are desirable (i.e., public health) over other things in the world that are desirable (i.e., short-term economic interests). However, it could also be fruitful to explore how these judgements alter ideas in our minds about desirable things (Korf 2022). I thank Drew Schroeder for highlighting this ambiguity for me.

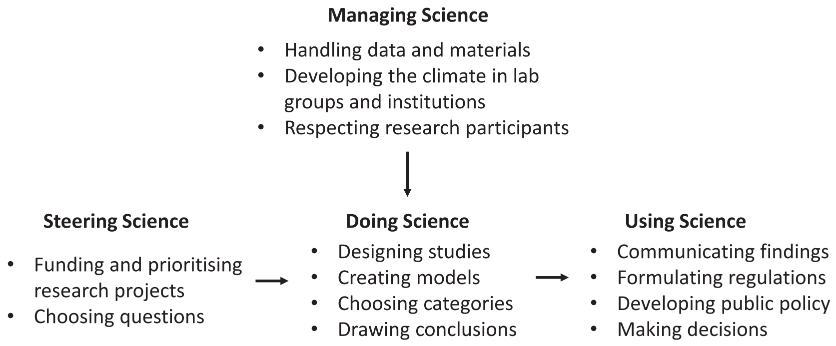

Figure 1 A representation of major ways in which values can relate to science

organises these judgements into four categories as a way of thinking about them in an organised fashion, but these categories are not intended to be mutually exclusive or exhaustive. Consider first the roles that values can play in steering research. The Introduction to this Element highlighted some of the ways that values steered research in response to the COVID-19 pandemic. For example, many scientists were motivated to shift their lines of research in an effort to help alleviate the suffering caused by the pandemic. For similar reasons (and also presumably because of economic motivations), policymakers and corporate leaders unleashed a flood of government spending, public–private partnerships, and private investment designed to support scientific research on COVID-19. Thus, the pandemic illustrates that values can steer research at both the individual and the institutional levels.

The ways that values steered research in the COVID-19 case appear to have been largely positive, but there are other cases in which values have steered research in very negative ways. For example, scientists throughout history have been motivated by sexist and racist values to search for biological traits that could explain the alleged inferiority of women or of marginalised racial groups (see, e.g., Kourany 2020). In addition, the desire to protect national security has stimulated research on horrific methods of biological and chemical warfare that are difficult to justify from an ethical perspective (Barras and Greub 2014). And negative value influences need not always be so obvious. There has been a great deal of discussion about the ways disciplinary boundaries can constrain scientists' research agendas in ways that hinder the solution of grand social challenges like alleviating poverty, mitigating climate change, preventing and treating disease, and addressing food insecurity (Frodeman et al. 2017; Kreber 2009; Weingart and Padberg 2014). One can frame this as a situation in which judgements about how research should be organised are influenced by values (in this case, the valuing of disciplines) that have been embedded in the structure of institutions like universities and funding agencies.

It is also worth emphasising that the judgements involved in steering research involve more than just decisions about what topics to study; they can also involve more subtle judgements about what questions to ask and how to investigate them. For example, Hugh Lacey (1999) has argued that scientists can study the same research domain using very different research strategies, meaning that they focus on different kinds of data and develop different sorts of theories. He illustrates this claim using the field of agriculture, where it is obvious that different strategies tend to support different values. Lacey points out that the dominant agricultural research strategy has focused heavily on questions about how to maximise crop yields by manipulating the genetics of seeds and by promoting fertilisers and pesticides. This is a powerful strategy that has increased crop yields significantly, but in some cases, it has had harmful effects on rural communities and the environment. Lacey points out that one could instead ask broader questions about how to design agricultural systems that alleviate rural poverty and that are environmentally sustainable. These questions could lead to alternative research strategies that would presumably incorporate greater input from the social sciences and from ecology. Lacey's analysis of research strategies and the values associated with them is applicable to many different fields. For example, the field of toxicology currently focuses heavily on identifying potential toxic effects of industrial chemicals. One could instead shift towards a focus on collaborations between toxicologists and 'green' chemists in an effort to design safer chemicals (DeVito 2016).

These questions about what research questions to ask can begin to blur into the second category of relationships shown in Figure 1: values associated with *doing* research.[4] This category involves the design of studies, the analysis of data, and the interpretation of results. Sometimes, philosophers of science have called this the 'heart' of science because it focuses on the process of drawing conclusions from evidence (see, e.g., Douglas 2009). This is the part of science

[4] Stephanie Harvard and Eric Winsberg (2021) provide an illuminating discussion about why the activities involved in steering science blur into the activities involved in doing science. They point out that scientific reasoning incorporates at least two activities: (1) representing phenomena and (2) drawing inferences. The activities involved in representing phenomena (e.g., creating models) straddle the line between steering science and doing science. On one hand, these representational activities involve choices about what phenomena to study and what questions to ask about them (i.e., steering science). On the other hand, these representational activities also have the potential to influence the conclusions that scientists ultimately draw (i.e., doing science; Elliott and McKaughan 2009; Okruhlik 1994; Winsberg 2018, 137). Harvard and Winsberg (2021) point out that the literature on values in science has been particularly focused on the judgements involved in drawing inferences, but it would benefit by placing more attention on the roles that value-laden judgements play in choosing and employing scientific representations.

where people have been most likely to think that values must be excluded. As discussed earlier, if science is supposed to be focused on obtaining accurate information about the world, then values – at least non-epistemic values – seem to be irrelevant to achieving that goal. Once one begins to consider the complexity of actual scientific research and the array of different ways that values can relate to this part of science, however, this view becomes more difficult to maintain.

Consider first the design of research studies. In an influential examination of previous studies on divorce, Elizabeth Anderson (2004) showed how researchers' values can influence how they conceptualise phenomena, which in turn influences how they design their studies. For example, she points out that earlier studies of divorce tended to conceptualise it as primarily negative, whereas more recent feminist research conceptualised it in a more open-ended way, as having the potential for both positive and negative consequences. Anderson showed that these different ways of conceptualising the phenomenon under investigation could affect how the researchers collected and analysed their data, thereby ultimately affecting their conclusions. To take another example, many environmental health scientists are now engaging in community-based participatory research (CBPR) projects. These projects typically involve collaborations between academic researchers and community members who have been affected by environmental pollution (see, e.g., Claudio 2000). Researchers who work on these community-based research projects often affirm that their involvement with community members affects how they design and interpret their studies, thereby ultimately making their work more responsive to the values and concerns of the affected community (Elliott 2017).

Another important aspect of doing research is the creation of models, such as the models used for predicting the impacts of climate change. Philosophers of science have identified a number of ways in which values can intersect with the creation of these models (e.g., Biddle and Winsberg 2010; Harvard and Winsberg 2021; Schienke et al. 2011). For example, Kristen Intemann (2015) has pointed out that climate modellers often have to make value-laden judgements about how to optimise their models; for instance, it might not be possible to optimise a model to predict both local changes in precipitation and global temperature changes, so scientists might have to choose between those two goals. Similarly, Wendy Parker and Greg Lusk (2019) have shown that when climate service organisations use models to make predictions about the future threats faced by local communities, they have to make numerous value-laden judgements about which models to use and how to combine information from different models. For example, if these choices are made in ways that generate