



Delft University of Technology

The Guidelines Challenge

Copeland, S.M.

DOI

[10.1007/978-3-030-41239-5_6](https://doi.org/10.1007/978-3-030-41239-5_6)

Publication date

2020

Document Version

Final published version

Published in

Rethinking Causality, Complexity and Evidence for the Unique Patient

Citation (APA)

Copeland, S. M. (2020). The Guidelines Challenge. In *Rethinking Causality, Complexity and Evidence for the Unique Patient: A CauseHealth Resource for Healthcare Professionals and the Clinical Encounter* (pp. 95-110). Springer. https://doi.org/10.1007/978-3-030-41239-5_6

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Chapter 6

The Guidelines Challenge



Samantha Copeland

6.1 The Tension Within

6.1.1 Evidence Based Medicine and the Rise of Guidelines

Part of the reason behind the popularity of evidence-based medicine, or EBM, in the last few decades has been its promise to standardise care. When Guyatt et al. (1992) introduced the concept and practice of EBM in medical education, it was in direct response to the ‘old ways’ of medicine. As they depict it, medicine was a profession where well-known doctors were looked to as the authorities about healthcare: expertise came from experience, and was demonstrated by the admiration of one’s peers. Standards of care could vary widely from hospital to hospital, and new evidence about what practices and interventions worked best was rarely taken up by institutions (Guyatt et al. 1992). The appeal to the ‘best evidence’ as the driver of the standard of care meant that everyone could have access to the reasons behind why things were done in a certain way; such reasons ought to be objectively good reasons, rather than merely authoritative. Patients could expect the same quality of care from every practitioner, and doctors would know the best thing to do for their patients.

For many reasons, this ideal was hard to realise in the world. First, because individual doctors could not possibly be expected to study all the available evidence for any given intervention, and so assessments of what is the best evidence and the recommended strategy for care had to be centralized. Organizations like Cochrane

In this chapter, we look at guidelines as an example of a medical institution and practice that would be affected by a shift in ontology from the current biomedical models toward a dispositional model of causality.

S. Copeland (✉)
Ethics and Philosophy of Technology Section, Delft University of Technology,
Delft, The Netherlands
e-mail: s.m.copeland@tudelft.nl

gained in authority, becoming the sources of ‘systematic reviews’ that did this work on behalf of practitioners. Second, because standardizing care takes more than everyone having access to the best evidence about which practices are optimal. It takes governance, and the organization of institutions with many levels of care and expertise, with different needs and different jobs. Furthermore, medical institutions, in many nations, are public institutions. Consequently, their governance is the concern of national governments, too, and standards must be set for ethical and political reasons. As a result of these and other influences, healthcare guidelines became a resource for governments, medical institutions, and practitioners alike. They are a means for standardizing care, they are meant to communicate which practices and decisions the best evidence says is optimal, and they are both a guide to uncertain practitioners and the means to inform patients of their options.

6.1.2 Guidelines in Practice

However, given what has been said in this book so far, it seems there must be a conflict between standardizing care, as guidelines seek to do, and using the best available evidence to provide optimal care for the individual patient, as practitioners who use guidelines seek to do. That is, the possibilities offered by dispositional causality for understanding health, disease, and the effectiveness (or not) of medical interventions, point to a tension in the goals of guidelines, under the EBM interpretation of what they can and should do.

This tension is found at a larger scale as well. That is, governments take a public health approach to medicine, and for that purpose, statistics are indeed ideal. Clinicians on the ground in medical care, however, find themselves dealing with individual patients in unique situations. There is not always a box to tick to describe their patient, or what they have done to care for them, on the standard forms they are required to fill out. Insofar as EBM encourages the standardization of care and the quantification of medical evidence, it caters to the public management approach to healthcare that is currently frustrating many healthcare practitioners.

In this chapter, we take a look at this tension, and give some ideas about how we can use the philosophical tools introduced so far to resolve it. Specifically, we focus on how this tension has manifested itself in the debate about healthcare guidelines. Healthcare guidelines in one sense exemplify the standardization of medical care. They present a predetermined list of options, they identify and plan a treatment course for a patient according to the category that fits them best, and they are generated by interdisciplinary committees using the best available evidence. However, they are simultaneously an honest attempt to provide guidance to the practicing, individual clinician, who hasn’t the time to delve into all the relevant evidence herself, and must rely on others to assemble the best available evidence into a format she can use in her daily, busy practice. Thus, the guideline can be both a hegemonic force, representing consensus and conformity, and also a helpful aide when a clinician is truly in need of information and guidance—they are meant to be both rule

and resource. This dual role creates tension in the clinical encounter because following the rule is not always the best thing to do in a particular situation, and so in many cases what the clinician needs is a resource that is not at all a rule. Guidelines, then, to truly aide the clinician, must be something other than rules.

These tensions can be truly problematic, when clinicians are asked or are otherwise inclined to believe it is better to follow guidelines even when their expert judgment would suggest a different path.

What singular question could be more pressing for clinicians today: how do we prepare the way for the return of the person in contemporary healthcare amidst rife healthcare commodification and the mechanical one-size-fits-all approach that is EBM? ...In clinics and hospitals the world-over, the narrative, the personal and the biographical person is often found bullied and threatened by this hegemony that is contemporary EBM; healthcare is de-personalised.

Chris Worsfold, 'Learn to stop worrying and love evidence based medicine', CauseHealth blog <https://causehealthblog.wordpress.com/2016/02/07>

6.2 Guidelines and Tramlines

A point that was frequently brought up in discussions during CauseHealth events was the worry that guidelines were often regarded as rules to be followed, rather than as guidance about what care options are available and appropriate. This phenomenon is captured by the frequently used phrase 'Guidelines not tramlines'—rather than a set track, the following of which is not a choice but rather given, guidelines ought to offer choices to clinicians and patients. The tendency to treat guidelines as though they were tramlines is, as suggested above, part and parcel of the public management approach to medicine, and presents clinicians with a conflict of interest: do they risk reprimand and even litigation, if they follow their expert clinical judgment to act in a way not advised by the guideline? Osteopath Stephen Tyreman thought this was a problem in health care, and promoted Values Based Practice (VBP) as a remedy:

VBP is being given further support by the recent [Montgomery judgement](#) [UK] in which 'doing the right thing' for a patient is not just following the guidelines or taking account of risks as a percentage, but taking into account 'what matters' to the patient, what his/her values are. This involves a different kind of decision-making and can't be achieved by reading out all the risks, or ticking a box that says we've told the patient the implications of the procedure being recommended. It involves asking how this treatment might impact on the patient's life and what matters to them. This requires a proper dialogue. What the regulators realise they are up against in supporting this is that practitioners will not be able to follow a safe procedure—an algorithm or guideline—that will guarantee they do the right thing; they must make a judgment weighing up a range of fact and value-based criteria, which of course is what David Sackett originally intended EBM to be (e.g. Sackett et al. 1996). The response from some practitioners is to say, 'we haven't got time to do all this deep and meaningful discussion stuff, just tell me what the right thing to do is.' Or, 'what if I make the wrong decision? I don't want to get into trouble.'

Stephen Tyreman, 'Standards for regulation', CauseHealth blog <https://causehealthblog.wordpress.com/2016/12/15/standards-for-regulation/>

6.2.1 Guidelines and Evidence Based Policy

One of the issues is the content of guidelines themselves. When designed, for instance, to effect changes in the management of healthcare, guidelines indeed present as rules to be followed. Unless everyone follows such guidelines, the changes won't happen. For example, there may be a guideline about how much time to spend with a patient on average. Or, there may be specific guidelines about what steps to take with patients who have certain diagnoses—which tests ought to be done, which specialists engaged, or even how long a session of cognitive behavioural therapy is allowed. When such guidelines are tied to billing processes in the healthcare system, they can become constraints on what is possible for clinicians to do or to offer their patients, rather than guidance about what might help. One could argue, however, that guidelines that focus on management issues really ought to be treated as tramlines, since they require maximum compliance in order to be effective. As an example, expensive experimental therapies should not be used up on patients who cannot clearly benefit from them, for instance because they are in a terminal phase of life. It is hard to think of a good reason to deviate from this managerial guideline. But in the example above, where the length of a cognitive therapy session is mandated rather than suggested, management-focused guidelines seem to cross over into the territory of the clinical encounter.

Another example of this was given at the Guidelines Challenge Cause Health conference by speaker Hálfdán Pétursson (Pétursson et al. 2009; for a review of talks at that conference, see Anjum et al. 2018). He discussed the problems encountered by general practitioners (GPs) who are trying to follow the guidelines for preventing cardiovascular disease in Norway. If the guidelines are followed as written, it would require more hours of work than the entire GP workforce in the country could put in, just for the one task of preventing cardiovascular disease in those who are deemed at risk. As this example shows, not all aspects of healthcare are manageable through rules, and sometimes implementing the best evidence in an ideal way is simply impractical.

Further, when we bring what has been said so far in this book into play in this discussion of the problem of guidelines, we can see that there are also ontological and epistemological reasons behind the problematic but common view that guidelines present tramline-type rules rather than mere guidance. We explore these reasons in the next two sections, as well as how the philosophical framework of dispositionalism might offer a resolution to these guidelines-related challenges.

6.3 The Ontology of Guidelines

The ontology of guidelines can mean two different things. It can mean the nature of a guideline itself—its function, its form, and what goes into developing it—or it could point to the ontological assumptions that shape guidelines as they are

developed and used. It is the position of CauseHealth that these two meanings of ontology are intertwined. How we think guidelines ought to be used reflects, and in turn is reflected in, what we think guidelines are about, and why we think they can be effective.

One might say that healthcare guidelines exemplify a utilitarian approach, for instance. Utilitarians like Jeremy Bentham argue that the best (moral, correct) action is the one that will result in the greatest utility (pleasure, or happiness) for the greatest number of people overall. So, a frequentist utilitarian would argue that guidelines should be designed as rules to follow that, overall and if everyone follows them, are most likely to lead to the best possible outcome for the highest number of people. This works well at a population level: consider, for instance, the importance of ‘herd immunity’ as a reason for making a vaccination mandatory for all healthy members of a population. But this, as we have seen, is not the same as obtaining the best outcome for a particular person at a particular time, which is generally the goal in the clinical encounter.

Simply put...

Utilitarianism is a form of consequentialism, the ethical framework that assesses moral action according to its impact on others. That is, we need to consider the consequences of the action, to know if it is the right one, when compared with the consequences of other possible actions. Bentham popularized the idea of ‘utility’ as a way of measuring and comparing the impact of possible actions—actions with greater utility caused more pleasure and less pain overall.

The nature of guidelines is just what we discussed in the last section: are guidelines rules to be followed, or are they collections of good advice and a presentation of relevant options to a clinician and her patient? In order not to treat guidance for the clinical encounter as a managerial rule, we have said that it is important to avoid seeing guidelines as tramlines. But more than this, it is also important to conceive of the function, form, and creation of a guideline in the right way, so that they are developed and used correctly and effectively. The difficulty is in resolving the tension between the need for flexibility, to allow for the particularities of the clinical encounter to influence decisions about care, and the need for standardization of access for all patients to quality care.

6.3.1 Logically Speaking, Guidelines Cannot Be Rules

Anjum and Mumford have written on the nature of guidelines (2017), arguing that guidelines must, logically speaking, be mere guidance rather than hard rules. Even if a guideline is effective as a rule, it still doesn’t help the clinician decide what to

do in the individual case. That is, the clinician must, in each case, still weigh up whether the guideline offers the best path or options for treatment for the particular patient in care. Thus, the guideline cannot be designed as a rule that can be universally followed, given a diagnosis or situation; because its relevance can be questioned in reflection upon the particularities of an individual patient, the guideline itself offers a choice, rather than a rule. For instance, in Chap. 4 of this book, Anjum and Rocca describe the case of guidelines for morbid obesity. While current guidelines do offer choices to the clinician, all of those choices rely on biological conceptions of obesity only. When practitioners encounter a case of morbid obesity that seems to be caused by trauma instead, then they must choose to go outside the guidelines. Kai Brynjer Hagen (see Chap. 10, this book) suggests that the best available evidence in such cases is not captured in guidelines, but rather in the patients' narratives themselves. In cases like these, a clinician must himself evaluate what the best evidence is, in choosing whether to follow the guideline at all.

That is, the complexity of a patient's situation seems to call for a unique approach to their treatment, rather than the application of a general rule. Thus, the ontology of the patient requires that guidelines be treated as advisory rather than regulatory, as a resource rather than a rule. This then calls into question the idea of guidelines as being rule-utilitarian—there will be no ‘rule’ that benefits more than the singular patient to whom it will absolutely apply (presuming such a patient does in fact exist, which is up for debate if guidelines are written for statistically ‘average’ patients, for instance, see Anjum, Chap. 2, this book). Even when a guideline represents the treatment paths that tend to work, they will not always present the best thing to do in a particular case. For this, judgement is needed, and familiarity with the patient himself.

Simply put...

Rule-utilitarianism suggests that we can find a rule that, if everyone follows it, will provide the most utility for the most people, overall. Act-utilitarianism suggests, in contrast, that each individual decision about whether to take a moral action must weigh up the potential consequences and utility of that particular action.

Therefore, rather than rule-utilitarianism, we might accept that act-utilitarianism is the correct way to understand the nature of guidelines: act-utilitarianism allows one to consider the consequences for overall utility of a specific act or decision, rather than of a generally applied rule (for a full review of this argument, see Anjum and Mumford 2017). However, note that this moves the source of the guideline’s utility from the guideline itself to the decision about whether to use the guideline: it is not about the nature of guidelines anymore, if we follow this line of reasoning, but about the nature of the clinician’s decision in that particular case. Consequently, it is not guidelines but the clinician who would be utilitarian in nature, after all.

There is thus an ontological reason for *not* seeing guidelines as tramlines, even if we do consider them to be utilitarian in nature. We cannot think of guidelines, ontologically, as utilitarian rules that represent (from a public health perspective) the best treatment options for the greatest number of people. In order to bring the most utility to the most people, rather, each case must be considered individually, before it can be decided whether the rule should apply. So the rule is not utilitarian on its own. Rather, it is the act of using or not using the rule that could be utilitarian, if making that decision for that patient also brings about the most possible utility overall. It seems, then, that the best action to take will be the action that brings about the best consequences for that particular patient. Even when we look at it in terms of the most utility for the most patients over all, this will happen only when each individual patient is given the best possible care for that patient specifically (a singularist approach) rather than when all patients who fit into a generalized category are all given the same care (a frequentist approach). But how do we make guidelines specific enough to give the best care to each individual patient, when patients can differ greatly in practice? Before addressing this problem, we will explore the practical reasons for making guidelines more specific than general.

6.3.2 What Does This Mean for Guidelines in Practice?

There are practical reasons for not seeing guidelines as tramlines as well; practitioners have good reasons for resisting the imposition of guidelines as rules for their practice. As we saw above, pressure to follow guidelines, and the presumption that they capture the best available evidence, means that practitioners may fear repercussions for not treating them as rules. But, as we see in the obesity case, guidelines do not always capture the best available evidence in relation to the particular patient at hand. So, when we take dispositional causality and person centered medicine as our paradigms for medical science and care, we can see there is a serious conflict between creating a rule to follow in the clinical encounter and using the best evidence available to decide on a course of care. It is simply not possible for a guideline to do both.

Indeed, one of the best known explorations of this problem has been written up by Gabbay and le May (2004), in their introduction to the idea of seeing guidelines as ‘mindlines’ instead of tramlines. Gabbay and le May show that guidelines play a complementary role to other practices for healthcare practitioners, including referring to other known authorities in their professional networks. That is, guidelines act as additional sources of evidence about best practices, weighed up in relation to practices they already rely upon in deciding what the best available treatment options for their patients are. Practitioners use their own judgment to decide when and if they will incorporate the advice given in guidelines into their practice and decisions. So guidelines are not treated as rules to follow in practice; thus, it does not make sense to develop them or to try to enforce them as such rules.

The best evidence available is going to include evidence related to the particularities of a situation—to the dispositions of the particular patient, and to the specific context of that patient as a person. But this kind of evidence cannot support the formulation of a rule that should be followed beyond this particular case. Once we understand that tramlines simply cannot do the work that guidelines are supposed to do, any conception of guidelines that insists they be followed as rules is wrongly conceived.

6.4 The Epistemology of Guidelines

Guidelines are meant to represent the best options for care in a given situation, given the best evidence available. But this recalls the problem of how we can know which of the available evidence is, in fact, the best. Guidelines developers regularly struggle with this problem. Many of them have adopted a system called ‘GRADE’ing the evidence; the principles of the GRADE (the Grading of Recommendations Assessment, Development and Evaluation) working group can be followed in order to rank evidence from best to least, in respect to how reliable it is for grounding guidelines. The claim is that this methodology offers greater flexibility than the original hierarchies of evidence developed by EBM proponents, because it incorporates a wider scope of evidence and then evaluates that evidence via a greater number of parameters. Whereas the hierarchies relied upon the methods by which evidence was generated to rank one kind of evidence higher or lower than another, GRADE methodology takes into consideration perceptions about the reliability of specific research results, sometimes allowing, for example, evidence that was generated by a low-level method on the EBM hierarchy to be ranked higher, and vice versa. For example, if an observational study reported a large enough effect, confidence in its results would be strong enough to rank it alongside an RCT; or, if there were a risk of bias in the way an RCT had been conducted, for instance if not enough women had been included as participants in a drug study, then its results would be ranked with lower confidence, especially in respect to women. However, even this more flexible way of ranking evidence still limits itself to considering the best available evidence to come from organised, clinical trials. This leaves out, for example, mechanistic evidence, which many believe to be of highest importance for determining causality, and patient narratives, key to understanding causality in the single case (see Anjum and Rocca, Chap. 4, this book).

A second epistemological concern in the development and use of guidelines is how to integrate this wider scope of evidence into a single or set of decision-making options for clinicians and patients. Different kinds of evidence require different kinds of expertise and even different kinds of reasoning to be employed in their evaluation. This is the task taken up in recent years, for instance, by the GIN

(Guidelines International Network) AID (Appraising and Including Different) Knowledge working group (Wieringa et al. 2018). One way of tackling this is to engage a diverse group of experts in the developing of guidelines, who are able to evaluate a wider scope of evidence through their joint efforts (Zuiderent-Jerak et al. 2012). A second way is to improve upon the transparency of guidelines, including details about the way in which decisions were made by developers as well as the end results of their deliberations.

6.4.1 Transparency and the Tension Between Flexibility and Standardization

Transparency, that is, is one of the ways that more flexibility is being built into guidelines by developers. When developers are able to be transparent about the choices they have made—for instance, about the reasons for why they made the decisions they did about which evidence was best, and how that evidence translates into the best practices they describe in the guideline—then clinicians have better tools for deciding whether they ought to follow the guideline itself. A clinician might, for instance, disagree with the reasons behind the choices that developers have made, insofar as that clinician would not have made the same choices either generally speaking or in respect to the particular patient at hand.

Of course, there are limits to how much transparency is useful. Too much transparency could mean that guidelines offer no better guidance than the evidence itself to clinicians and patients, who must first decide whether they agree with the experts in order to make use of their guidelines. An effective guideline is developed by experts who can be trusted to evaluate the evidence on the behalf of others who lack either the expertise or the time at the moment when decisions must be made to do that evaluating work themselves. Again, this follows the observations of Gabbay and le May (2004), that practitioners are more likely to accept new evidence when it is promoted as good evidence by trusted colleagues and authorities in the field.

The ontological concerns raised in the previous chapters and section come into play here, as well. Knowing what evidence is best, and knowing what treatment options it is evidence for, requires an ontological judgment about evidence—we have to first have an idea of what good evidence really is, to tell if the evidence we have is also good. Thus, any process by which experts come together to evaluate evidence starts from their assumptions about what makes good evidence. In developing a guideline, it can be supposed that the best evidence available is also the best evidence for the purpose of developing a guideline. If we further see guidelines as rules to be followed, and a good outcome as achieving the best results most often for the most people, then we are evaluating evidence in terms of its quality as a support for a rule. As we say above, this backs us into an uncomfortable corner, in the thick of the tension between flexibility and standardization.

6.4.2 When Should the Particular Be Engaged?

The epistemological factor that needs to be highlighted here is the point at which nuance and particularities become relevant for whether a guideline will be useful. This point, I argue now, is whenever decisions must be made about what are the best available healthcare options for the particular patient at hand. For these decisions, evidence about the patient's condition and evidence about what treatment options are best must come together. Guidelines are meant to offer guidance on how to do this. Often, this point of convergence is assumed to be the point of diagnosis. But, as previous chapters have shown, there is much more to determining which treatments are best than diagnosing the patient. Further, in many cases, no accepted diagnosis fits (see Anjum and Rocca, Chap. 4, this book). And, diagnosis should not present a point when the clinician has to decide which guideline to follow, or whether to follow one at all. Guidelines should be ever-present tools for the clinician, no matter the diagnosis, leading them through the resources available in a way that really helps. So the particularities should be an inherent part of guidelines, acting for the clinician as starting points from which to begin a search, for example.

Additionally, guidance is needed from the moment the clinical encounter begins because, from the very beginning of the encounter, a clinician is observing new and varied evidence about her patient's health, and thus needs correlating knowledge about what treatment options might be available and best. Thus, guidelines ought to offer guidance not only after a diagnosis has been settled upon, but from the moment the clinical encounter begins.

And finally, we consider flexibility to be important in respect not only to how guidelines can be used, but also in how they might be formed. The clinical encounter, that is, should be seen as a resource for gathering further evidence, not only about the individual patient but also about the effectiveness of the guideline itself or the usefulness of the evidence available in respect to that patient and that clinical encounter. New knowledge is gathered at the site of care, and guidelines should not be imposed in a way that constrains such knowledge production, but rather can play a role in enabling it. For instance, as Rocca (2017) recommends, building bridges that enable evidence from the clinical encounter to be taken up by scientific researchers may be the best way to gather mechanistic evidence, which must be observed in the singular case and reported along with its qualitative context. Further, unexpected discoveries frequently occur within the context of the clinical encounter (such as unusual responses to drugs, whether adverse or beneficial), and processes for developing and using guidelines ought to not only allow for this to happen, but also be prepared to learn from serendipity—fortunate, though unexpected discoveries—when it happens (Rocca et al. 2019). Guidelines, in the right form, could build such bridges and enable serendipity.

All of this means that attempts to build in those particularities and nuance by developing any single guideline that will work as a linear and certain pathway through the steps of treatment are misdirected. No matter how wide the scope of evidence, or diverse the expertise and reasoning employed, guidelines developers cannot hope to capture all details that may be needed by a decision-making team of clinician and patient in the clinical encounter. Therefore, it cannot be the function of

guidelines to offer a complete set of evaluations and options. Rather, as suggested in Chap. 3, local knowledge, patient narratives and a variety of evidence are needed to determine the propensities that are relevant in this clinical situation (Rocca, Chap. 3, this book). That is, the epistemic role of guidelines cannot be to provide all the knowledge that will be needed in this context (i.e., all the material that will be needed to make the best possible decisions about care).

Let's return again to the idea of Gabbay and le May (2004), that 'mindlines' are created through the tacit and ubiquitous interactions between colleagues and experts within the practice of healthcare, in contrast to guidelines being imposed from outside and accepted as new rules to follow in and of themselves. Guidelines, in their depiction, are useful insofar as they are trusted sources of new evidence. To be trusted sources, they need to be trusted by trustworthy colleagues, for instance; similarly, if they are developed by trustworthy experts, it is the expertise and not the evidence that will be followed by practitioners. So any guideline will have to be transparent, at least, about who made the decisions when creating it, what their claims to expertise in the interpretation and application of such evidence are, and why they made the decisions they made. It is these decisions, then, that are the true content of a guideline (the guidance they offer), and not the actual treatment paths that are recommended. This gives us one way of understanding how guidelines can be formed and transmitted in a useful way—as collections of expert advice, rather than as a simplified rule or set of options.

Consider one of the biggest challenges for guidelines developers—the prevalence of comorbidity among healthcare patients. As we have discussed (see Rocca and Anjum, Chap. 5, this book), the tendency is to divide the medical body (both the body of the patient and the body of medicine's institutions) into disparate parts. But the reality is that many patients present with multiple conditions, which interact in ways that single-disease guidelines may fail to capture.

A related challenge is how to deal with the wealth of potentially relevant particularities, and how does one identify patterns (that is, decide what counts as evidence) in a unique situation? Guidelines must not only allow for intersection, but also for the flexibility that is needed in the face of the variety between patients and the fact that each patient will also change over time. They need to provide the means for telling a causal story about that patient, one that includes emerging causal relations and genuine complexity (see Rocca and Anjum, Chap. 5, this book).

So guidelines must be broad enough in scope to be useful before diagnosis as well as after, and they must be flexible enough to be adaptable to a changing and unique situation. The epistemic role of guidelines, then, is to offer a navigable network of connections between observable evidence in the clinical encounter and generalizable evidence obtained via the results of medical research, as well as information about what treatment options are actually available and likely to be beneficial. That is, rather than a flowchart of pathways to take in a certain direction once the starting point is decided upon, a guideline could present more like a web of data and information that can be searched in any direction and that would highlight interconnections between possible pathways. These pathways ought to come in the form of expert advice on how to integrate the best available evidence into practices that are already in play. Rather than an imposition of a new rule upon practice, then,

guidelines would act as a resource for ways to improve upon existing practice, allowing for variation between practices and for the best available evidence to change practice from the bottom up, rather than the top down. This is just one way to imagine how guidelines might work, if we start from a dispositionalist perspective.

To suggest that guidelines should work bottom up, should trace the interaction between expertise and the best available evidence in a way that is both transparent and helpful to practitioners, and should allow for variations within practice and between institutions, however, seems to bring us back to the original problem identified by the proponents of EBM—that is, how do we standardize best practices in healthcare without also linking best practices to authority figures instead of to the evidence? We have argued that the solution is not to create guidelines in a way that makes them rules to follow, or as inflexible as tramlines. In the next section, we explore further how dispositionalism allows us to imagine a middle way, by linking best evidence directly to the individual, and promoting singularism over frequentist or utilitarian approaches to the standardization of best practices.

6.5 Guidelines in the Dispositionalist Way

In moral philosophy, the idea that right and wrong can be defined by a system of rules has been challenged. If nothing else, it's obvious that two rules could easily come into conflict and one of them has to be sacrificed. Telling the truth is good but not necessarily if a killer asks for the whereabouts of any intended victim. There could be circumstances in which it is right to lie. In response, Dancy [2004] proposed a theory of moral particularism, a view in which each situation has to be understood as complex and requiring its own moral assessment, which could well be unique and unrepeatable. I would favour coupling this with a strongly dispositional version of virtue ethics. Telling the truth tends to be right but not necessarily so. Assessing the whole complex of circumstances might weigh in favour of lying.

Now I think the issue of how to understand and use a guideline clearly relates to this discussion. We could interpret a guideline in a dispositional way rather than as an absolute rule. A particular intervention may tend to relieve a particular symptom but in many contexts it need not be the right intervention to prescribe. If this is right, I think it would be to the benefit of all stakeholders—clinicians, guideline bodies, regulatory authorities, and patients—to understand dispositionalism and particularism. This could be a challenge when rule-based laws and codes of ethics are easy to grasp, but there is a potential benefit to be gained from pushing ahead for a conceptual change.

Stephen Mumford, ‘The Notion of Guidelines’, CauseHealth blog
<https://causehealthblog.wordpress.com/2016/12/08/the-notion-of-guideline/>

6.5.1 So, What Should We Do with Guidelines?

We wrote earlier about the fact that we cannot know about a causal relationship until after it has been observed (see Anjum, Chap. 2, this book). However, even after we have seen the same cause and same effect occur together repeatedly, this does not

mean that the causal relationship did not, in that first instance, actually exist. That is, our epistemological state does not determine the ontological case. In the case of guidelines, it seems we are trying to use what we know, epistemically speaking, to say something about what must be true, ontologically speaking. When we set a guideline, we pass a judgment about what treatment options are best, given our evaluation of the evidence available. This judgment is not just a guess, it came about via rational deliberation among experts. However, the best evidence available, as has been shown in previous chapters, still cannot tell us what will actually happen, in a particular case, given the specificity of the mutual manifestation partners involved. In most cases, the best available evidence can only give us a probability range of an outcome occurring; epistemically speaking, probability is often the best knowledge we can have. Thus, guidelines are still useful, even if they cannot prescribe exactly what the best decision will be.

But dispositionalism gives us something further to work with than this. It is true that we cannot use our knowledge to determine exactly what will happen when a particular patient embarks on a particular treatment path. However, as I mentioned in the last section, the clinical encounter itself provides considerable observable evidence about the particularities and nuance needed in deciding what options presented in a guideline to follow. What is important is that guidelines be written in such a way that the general medical knowledge we have already can be met effectively by the knowledge gained during the clinical encounter itself. What the clinician might know to be true about a situation gives her new tools for evaluating the relevance of a guideline for her patient.

What might a guideline look like, then, if we take up the conclusions drawn in this book? Above, we said that they ought to offer a navigable network of connections between observable evidence in the clinical encounter and generalizable evidence obtained via the results of medical research, as well as information about what treatment options are actually available and likely to be beneficial. Mumford, in the quotation above, adds the concerns of particularism and virtue to be considered by those who develop and use guidelines (see also Anjum and Mumford 2017). Add to this what I have brought up in this section, that guidelines provide only half of what is needed for their own effectiveness in the clinical encounter—it is up to the clinician to make observations and the patient to contribute further evidence in order for them to work successfully with the guideline to make decisions. Taking all this together, it could be argued that guidelines are most effective as tools for use in the clinical encounter when they are transparent and accessible in a way that allows patients and clinicians to question the guideline itself where necessary, and to easily draw connections between their particular situation and the more generalizable advice given by the guideline.

In order to be fully dispositional, however, guidelines also need to provide more information about what kinds of mechanisms might be in play, and what kinds of dispositions are likely to affect the treatment path's effectiveness, if chosen. This will require not only a transparent process, but an iterative process, that takes up new evidence gained in clinical encounters and makes it accessible to future users of those guidelines. As in pharmacovigilance, patient reports can play a key role

here (Rocca et al. 2019). With the increasing use of data mining techniques and electronic records, paired with clinician and patient narratives when helpful, such information is already taken up in guideline development—it could, perhaps, more prominently shape the way guidelines work, so that using guidelines also means actively engaging with them on a regular basis. The experts whose advice will be accepted by practitioners, as we have seen, comes not only from guidelines developers, but also from colleagues and leaders in the field who have put them into practice and reflected upon their impact. So the development of guidelines (especially if we take them seriously as ‘mindlines’) does not stop at their implementation, but rather needs to be continued afterward by taking up and disseminating the ways in which individual practitioners have been able to use them in practice. In addition, this kind of new evidence can communicate details about what kinds of patients the guidelines have worked best for, and the contexts in which they have so worked.

In addition, guidelines could offer advice on how to interpret different kinds of evidence that will arise in the clinical encounter, such as what kinds of things to note within a patient narrative that may be clues to relevant dispositions, or other symptoms not immediately evident from a patient’s physiology but which may affect treatment options. In this way, guidelines could do more than offer a dispositionalist clinician the right guide to assessing and treating a patient, they could promote a dispositionalist approach to care. Rather than seeing guidelines as presenting ready-made options for a suitably assessed patient, then, guidelines would provide suggestions for how to assess the situation and how to take various psychological and social factors into account when making decisions with patients about their care. Many of them, indeed, are already taking this task on.

Such an approach may seem complicated, but only if we hold on to the utilitarian ideal of finding guidelines that maximise the utility of medical care over a population. Once we have taken a dispositionalist turn, toward the particular patient and the integration of a plurality of evidence types in order to focus on mutual manifestation potentials, then it makes no sense to spend our time developing such rules, if they are not also useful. Rather, it would be more natural to create guidelines that work more like networks and databases, giving access to a constantly updating base of evidence and resulting advice, and highlighting the fact that evidence collection is a continuous process, not ending at the point of diagnosis and treatment choice.

Finally, we at CauseHealth recommend taking an ecological turn in medicine (see Rocca and Anjum, Chap. 5, this book). Guidelines could be part of such a major change in focus, away from specialist focus on different parts or systems of the body, and toward integration with extra-physiological factors. Just as recognizing mutual manifestations, complexity and causal singularism also means that we cannot best study a causal relation in isolation, we cannot expect to fully treat the so-caused condition in a patient by isolating it from its context and environment (see Price, Chap. 7, and Low, Chap. 8, this book). Similarly, it is unlikely that a single

guideline or set will fully treat that condition without also engaging the interacting factors. Consequently, and importantly, guidelines need to express an acceptance that many clinical decisions will be made in a state of uncertainty. We are unlikely to know, even with the best evidence available, whether a treatment plan will work in the predicted way, because it is unlikely that we will have full knowledge of all the mechanisms, dispositions and, thus, causal relations that may interact with that treatment in producing its effect in that particular patient. The ecological account, that is, takes seriously the complexity of the kinds of interactions that influence an individual's health. Once we take the ecological nature of health and healthcare into account, then, we see that the design of guidelines must also be a holistic process, seeing guidelines as sets of intersecting advice and increasing awareness of possible interactions—guidelines must be living resources, embodying both the uncertainty and the expertise of the clinicians and patients who use them.

6.6 To Sum Up...

One of the key problems in medicine today, as the clinicians we have had the pleasure to work with have told us, is how to handle guidelines. Public management approaches to medicine tend to promote guidelines as rules to follow, and clinicians often feel pressure to follow a guideline even when their judgment cautions them to do otherwise. This ‘tramline’ approach to guidelines, we have shown, is philosophically as well as practically problematic. Especially when we take dispositions as the ontology of the causal relations that guidelines want to key in on—the best way to cause a recovery, or to counteract the causes of a condition—we see that guidelines cannot and ought not be treated as rules to be followed.

If this is taken seriously, then the development of guidelines must be more than the collective effort of diverse experts on a particular category of disease or subgroup of patients. Rather, guidelines must give clinicians the tools to assess the potential for mutual manifestations between the patient and the treatment options. They must allow clinicians to be flexible with the treatment plan, making changes as the treatment progresses and new evidence arises. And finally, their development does not end with the creation of a rule, but rather continues with the collection of that new evidence, taken back up into the guidelines to provide a continuously improving resource for each new clinical encounter. These features are needed for guidelines to be useful resources for clinicians, and we have used dispositionalism to ground them in the very nature of causation. This chapter, then, has presented one way in which the ontological assumptions we hold about medicine directly affect how our medical institutions work.

References and Further Readings

- Anjum RL, Mumford S (2017) A philosophical argument against evidence-based policy. *J Eval Clin Pract* 23:1045–1050
- Anjum RL, Copeland S, Kerry R, Rocca E (2018) The guidelines challenge—philosophy, practice, policy. *J Eval Clin Pract* 24:1120–1126
- Dancy J (2004) Ethics without principles. Oxford University Press, Oxford
- Gabbay J, le May A (2004) Evidence based guidelines or collectively constructed “mindlines?” Ethnographic study of knowledge management in primary care. *BMJ* 329:1–5
- Guyatt G, Cairns JA, Churchill D et al (1992) Evidence-based medicine. A new approach to teaching the practice of medicine. *JAMA* 268:2420–2425
- Petursson H, Getz L, Sirudsson JA, Hetlevik I (2009) Current European guidelines for management of arterial hypertension: are they adequate for use in primary care? Modelling study base on the Norwegian HUNT 2 population. *BMC Fam Pract* 10:70
- Rocca E (2017) Bridging the boundaries between scientists and clinicians. Mechanistic hypotheses and patient stories in risk assessment of drugs. *J Eval Clin Pract* 23:114–120
- Rocca E, Copeland S, Edwards IR (2019) Pharmacovigilance as scientific discovery: an argument for trans-disciplinarity. *Drug Saf* 42:1115–1124
- Sackett DO, Rosenberg WM, Gray JAM, Haynes RB, Richardson WS (1996) Evidence based medicine: what it is and what it isn’t. *BMJ* 312:71–72
- Wieringa S, Dreesens D, on behalf of the AID Knowledge Working Group of the Guidelines International Network et al (2018) Different knowledge, different styles of reasoning: a challenge for guideline development. *BMJ EBM* 23:87–91
- Zuiderent-Jerak T, Forland F, Macbeth F (2012) Guidelines should reflect all knowledge, not just clinical trials. *BMJ* 345:36702

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter’s Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter’s Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

