Cambridge Quarterly of Healthcare Ethics (2023), 32: 2, 163–173
doi:10.1017/S0963180122000263

CAMBRIDGE UNIVERSITY PRESS

ARTICLE

Cost-Effectiveness, Incompleteness, and Discrimination

Anders Herlitz

Institute for Futures Studies, Box 591, 101 31 Stockholm, Sweden
Email: anders.herlitz@iffs.se

Abstract

This paper argues that cost-effectiveness analysis in the healthcare sector introduces a discrimination risk that has thus far been underappreciated and outlines some approaches one can take toward this. It is argued that appropriate standards used in cost-effectiveness analysis in the healthcare sector fail to always fully determine an optimal option, which entails that cost-effectiveness analysis often leaves decision makers with large sets of permissible options. Larger sets of permissible options increase the role of decision makers' biases, whims, and prejudices, which means that the discrimination risk increases. Two ways of mitigating this are identified: tinkering with standards used in the cost-effectiveness analysis and outlining anti-discrimination guidelines for decision makers.

Keywords: QALY; discrimination; cost-effectiveness; healthcare rationing

Cost-Effectiveness, Incompleteness and Discrimination

Cost-effectiveness analysis is a powerful decision tool that can be used to improve healthcare policy as well as priority-setting, and distribution of scarce resources.[1] In brief, it is a decision tool that helps measure how much "bang for the buck" different options generate and if followed by decision makers, it ensures that they maximize the good impact of their choice. By measuring how much good they can get from a limited budget, decision makers who rely on cost-effectiveness analysis can maximize the goodness of their decisions, thereby avoiding wasting resources and choosing allocations that are suboptimal. In this paper, I will argue that the use of cost-effectiveness analysis in the healthcare sector risks leading to discrimination, which means, at the very least, that cost-effectiveness analysis should be complemented with anti-discrimination guidelines, something which is rarely recognized.

The paper is structured as follows. In the first section, I introduce cost-effectiveness analysis in more detail; in the second section, I argue that we have reason to accept that appropriately designed cost-effectiveness analysis for the healthcare sector is incomplete in the sense that it will fail to always fully determine a best option (an option that is at least as good as all alternatives). In the third section, I argue that this incompleteness introduces a risk of influence from discriminatory attitudes into the decision process, which might entail that certain salient social groups are advantaged and others disadvantaged. In the fourth, concluding section, I explore different ways of dealing with the discrimination risk.

Cost-Effectiveness and Covering Values

Although the terms "cost-effectiveness" and "cost-effectiveness analysis" sometimes are used synonymously with health maximization or maximization of the value of health,[2] I suggest that we should be more careful and think of cost-effectiveness in a more abstract and general way. I suggest that we think of cost-effectiveness analysis as the method that evaluates choice options in terms of how resource-effectively *a certain objective* can be pursued and leave it open whether this objective is to maximize some specific, unweighted value, such as health, or not. The objective might of course be health

maximization or maximization of the value of health, but it need not be. In other words, cost-effectiveness analysis ranks choice options in accordance with their $/objective ratio. The nuance I am emphasizing by suggesting that the objective need not be specified matters because the idea that one ought to maximize the value of health rules out potential distributional concerns such as inequality aversion, while the method of cost-effectiveness analysis seems perfectly compatible with distributional concerns, as evidenced by recent interest in incorporating so-called "equity weights" in cost-effectiveness analysis.[3] Both a ranking of choice options based on $/health and a ranking of options based on $/equity-weighted health are, I maintain, examples of cost-effectiveness analyses.

Cost-effectiveness analysis evaluates, *compares*, and ranks how effectively different options meet a given objective. This means that what one chooses as objective will have a decisive influence on the comparisons and rankings generated by the cost-effectiveness analysis. In abstract terms, the objective that one relies on when one engages in cost-effectiveness analysis reflects the answer to the question of *with respect to what* one makes the comparisons when one ranks options in terms of their cost-effectiveness. In the value- and decision-theoretical literature on comparisons and comparability problems, this is sometimes referred to as the "covering consideration" or "covering value."[4] Comparisons are triadic in the sense that one never simply compares x and y, but always compares x and y with respect to something, z. For instance, a comparison simpliciter of Paris and Marseille makes little sense, whereas comparisons of the cities with respect to which of the cities lies farthest to the South (Marseille) or is most populous (Paris) are perfectly intelligible. A covering consideration/value is that with respect to which one makes a comparison (in the examples above: latitude and population). In cost-effectiveness analysis, the objective one relies on is the covering value. We can say that cost-effectiveness analysis ranks options with respect to their $/covering value ratios.

To further understand the covering value that cost-effectiveness analysis in the healthcare sector relies on, it is helpful to look at the literature on distributive justice. Borrowing from this literature, we can say that the covering value used in cost-effectiveness analysis needs to rely on both a "currency" and a "pattern."[5] In the healthcare sector, the covering value is used to evaluate how the distribution of a certain good (expressed in a "currency") meets the standard of a certain "pattern" in the alternative outcomes that follow the options in one's set of alternative choices. If QALY (Quality-Adjusted Life-Years) maximization is the covering value, cost-effectiveness analysis ranks options with respect to how much (the pattern) QALY (the currency) they generate. If equity-weighted Disability-Adjusted Life Years (DALY) lost minimization is the covering value, cost-effectiveness analysis ranks options with respect to how much equity-weighted (the pattern) DALY (the currency) they avert. The choice of currency, thus, reflects an answer to the question of what kinds of benefits one focuses on and values, for example, QALY, opportunities, or resources.[6] The choice of pattern, by contrast, reflects an answer to the question of what one thinks is a desirable pattern of distribution of the benefits in question, for example, maximization, more equal distributions, or distributions that are better for the worse off.[7]

Plausible currencies that can be adopted when cost-effectiveness analysis is used by decision makers in the healthcare sector are summary measures of goods that one has somewhat reliable and available data on. The most obvious contenders are summary measures of health of different kinds, or summary measures of the value of health: QALY, DALY, and so on.[8] Importantly, what these summary measures have in common is that they encapsulate (dis-)values in different dimensions. All summary measures of health and all summary measures of the value of health encapsulate the idea that, for instance, physical pain is bad, that premature death is bad, and that certain disabilities such as not having the ability to walk are bad. Thereby, they rely on and make some trade-offs between different things that are valuable.

Plausible patterns that can be adopted and incorporated into a covering value for cost-effectiveness analysis in the healthcare sector include those found in the theories discussed by philosophers of distributive justice. Those promoting maximization of goods adopt thinking similar to utilitarianism.[9] Those promoting maximization of goods, in conjuncture with the idea that benefits to the worse off matter more, adopt thinking similar to prioritarianism.[10] Those promoting ensuring that everyone gets a certain amount of goods adopt thinking similar to sufficientarianism.[11] Those promoting more equal distributions of goods adopt thinking similar to egalitarianism.[12] Which of these patterns of

cost-effectiveness analysis one ought to rely on is a contested issue about which reasonable people can (and do) disagree.

To summarize, cost-effectiveness analysis in the healthcare sector can be understood as a ranking of choice options based on their $/covering value ratios. Choosing a covering value for cost-effectiveness analysis in the healthcare sector requires choosing a pattern and a currency. Plausible patterns can be found in the literature on distributive justice. In the healthcare sector, all plausible covering values must encapsulate currencies of good that ascribe (dis-)value to several different value dimensions. At the very least, they must ascribe value to longevity and improvements of health-related quality of life.

Incompleteness

Once one settles for what type of currency and pattern one finds plausible (e.g., maximize the value of health), questions can be raised about the properties of the covering value. I will focus on one particular such property: is the covering value *complete* or is it *incomplete*? A complete covering value is a covering value for which it is true that whenever one applies it to a comparison between two items that have some value according to it, it is true that one of the items is better than the other, or that the two items are equally good. For instance, it is always true that if one makes a comparison between two items that have some height, one of the items will be taller than the other, or they will be equally tall; the covering value, *height*, is complete. An incomplete covering value is a covering value that is not complete, which means that it is possible that neither of two items that have some value according to the covering value is better than the other, nor are they equally good. For instance, it might not be true that whenever one compares two artists (e.g., Mozart and Michelangelo) in terms of creativity, one is more creative than the other, but they are not equally creative either; the covering value, *creativity*, would in such a situation be incomplete. In other words, a covering value that admits of unconventional comparative relations (e.g., incomparability) is a covering value that is incomplete. In this section, I will argue that we ought to accept that all appropriate covering values in the context of cost-effectiveness analysis in the healthcare sector are incomplete.

Incompleteness is coupled with incomparability (when a covering value can establish that some items are incomparable, it is incomplete), but there are also other comparability problems that actualize a similar phenomenon.[13] If a covering value is vague (e.g., if one uses the covering value *baldness* to rank people with respect to baldness), it can sometimes be the case that the covering value cannot fully determine which of two items is better than the other or that they are equally good.[14] If there are more than three positive comparative relations (more positive comparative relations than the trichotomy: better than, worse than, equally as good as), it can sometimes be the case that a covering value determines that two items are related to each other with some non-conventional comparative relation, such as "parity."[15] In the last decades, value theorists have argued about which of these explanations best account for hard cases of comparisons.[16]

Since indeterminacy, parity, and incomparability actualize decision problems that are very similar, and since this is not the place to argue for a certain value-theoretical explanation of what phenomenon fundamentally leads to these decision problems, I will, in what follows, focus on whether appropriate covering values in the healthcare sector are incomplete*:

> *Incompleteness**: A covering value, C, is incomplete* if and only if it fails to fully determine a trichotomous ranking of all of the items that according to it have some value.

By trichotomous ranking, I mean a ranking of all options in terms of better than, worse than, equally as good as. The clause that indicates that we are only concerned with items that have some value according to the covering value is introduced to exclude non-comparability problems.[17] An apple and Beethoven's ninth symphony are non-comparable with respect to the covering value *good bicyclist*, for example, and we can expect all of the covering values used with cost-effectiveness analysis in the healthcare sector to similarly deem some items to be non-comparable, for example, the number 3 and the Andromeda

galaxy. Non-comparability is not particularly interesting in the context of evaluating cost-effectiveness analysis and its problems.

If a covering value, C, is incomplete* it is possible that it cannot be determined that either of two options that have some C, O_X, and O_Y, is more C than the other, but it cannot be determined that they are equally as C either. In such situations, if O_X and O_Y are the only available options, C will fail to guide a decision maker to an option that can be fully determined to be optimal (at least as good as all alternatives) because there simply is no such option in the choice set.

I will provide three reasons to expect that an appropriate covering value for cost-effectiveness analysis in the healthcare sector is incomplete*: intuition, the small improvement argument and disability-discrimination avoidance. First, consider the argument that options that decision makers in the healthcare sector face are sometimes good in very different ways. For instance, decision makers in the healthcare sector might have to choose between whether to invest newly available financial resources in a maternity ward or in a geriatric ward; they might have to choose between covering the cost of new and costly medicine that can benefit a small group of very badly off people or covering the cost of cheap medicine that can benefit a large group of relatively well off people; or they might have to choose between building an additional clinic in a city that already has a good healthcare system with queues that are sometimes long or to build a new clinic in a rural area where access to healthcare is worse, but where there are few people demanding healthcare. There are allocation alternatives that are so different so that it seems impossible to establish that one is better than the other, although they do not seem equally good either. To make the point in a different, more abstract way, decision makers in the healthcare sector might face choices where they need to decide between extending the life of 10 elderly people with 1 year each or slightly improving the health-related quality of life of 1,000 young people for the rest of their lives. Intuitively, there will be some situations of this kind, where different options are good in very different ways, where it will seem absurd and just wrong to claim that the options are equally good, but equally absurd and wrong to say that either of the options is better than the other. If it cannot be determined that one option is better than the other, or that the options are equally good, the covering value would be incomplete*.

It might be posited that some people actually do have strong and clear intuitions in cases like these; however, strong and clear intuitions in these cases are highly controversial. Consider, for instance, the intuition that the value of life trumps all other values so that it is always better to extend the life of someone than to improve the health of a large group of people no matter how small the life-extension is and no matter how many people one can help. Some people might have this intuition, but it is controversial, and it is hard to see a consensus developing. Similarly, some might have a strong and clear intuition that one should always prioritize the young over those who have already had a long and healthy life, but again, this is a highly controversial view, and it is hard to imagine a consensus on around it. The fact that some people have strong intuitions in these cases is not an argument against incompleteness*, since those who have strong intuitions have intuitions that run in different directions, and it is hard to see how they could come to an agreement. This is, in itself, an indication that the covering value is incomplete*.

Second, consider the small improvement argument, first presented (with a different purpose) by economist Leonard Savage.[18] In the literature, this argument has in recent decades become one of the strongest arguments in defense of incommensurability, which entails incompleteness*.[19] In its abstract form, it can be presented as follows:

Small improvement argument:

- Assume that the relevant covering value, C, is multidimensional such that there are at least two ways, p and r, in which an alternative can be more C, that is, $C = f(p, r)$.
- Assume that neither of the ways in which an alternative can be more C, p and r, trumps the other in the sense that if some alternative is more p or r than another, it is always more C.
- Assume that the set of possible alternatives consists of alternatives the values of which, with respect to C, are fine-grained variations of p and r.
- Given these assumptions, there will be two alternatives, X and Y, such that X is significantly more p than Y and Y is significantly more r than X, but neither is more C than the other.

- Would X, under these conditions, necessarily be more C than Y if one improved it ever so slightly? If not, X and Y are not equally as C which means that C is incomplete*.

The small improvement argument can be (and has been) used in a large variety of contexts in order to establish that covering values are incomplete* (or the possibility of incommensurability, parity, or vagueness). If C is artistic creativity, the small improvement argument can be invoked to support the idea that Mozart and Michelangelo must relate to each other in some non-conventional way, because it does not seem to change their relation to each other with respect to creativity if we discovered that Mozart had composed an addition aria.[20] Similarly, if C is understood as needs and a need is a function of the severity of illness and capacity to benefit, the small improvement argument can be used to argue that the principle of need fails to fully determine what one ought to do in all situations.[21]

In this paper, I want to suggest that when applied to summary measures of health or summary measures of the value of health, the small improvement argument gives us reason to accept that these should be incomplete*. In the previous section, I underlined that summary measures of health and summary measures of the value of health are multidimensional; at the very least, they must rely on valuations both of morbidity and longevity. Furthermore, it is obvious that neither of these dimensions trumps the other. Any plausible approach to a summary measure of health or its value will reject the idea that any amount of increase in length of life will be more valuable than any amount of reduction of morbidity. It is simply false that it is better to extend one life with 1 minute than to cure 10,000 people of paraplegia. And any plausible approach to a summary measure of health or its value will of course also reject the idea that any amount of reduction of morbidity will be more valuable than any amount of increase in length of life. It is obviously false that it is better to cure someone from hangnails than to save someone's life so that they can go on and live a healthy life for another 10 years. Finally, the set of possible alternatives (logically possible health interventions and policies) consists of alternatives the values of which, with respect to a plausible approach to a summary measure of health or its value, are fine-grained variations of morbidity and longevity impact. In other words, the assumptions for the small improvement argument are met in the context of summary measures of health and in the context of summary measures of the value of health.

To spell out the argument, let us read C as an attractive summary measure of health or its value, p as health-related quality of life and r as longevity. In other words, let us think of the covering value, C, used by a plausible cost-effectiveness analysis as a function of health-related quality of life and longevity, in line with the suggestion in the previous section. Since neither health-related quality of life nor longevity is categorically and lexicographically more important than the other, there will exist two outcomes that are very different with respect to health-related quality of life and longevity, X and Y, for which it will be false that either is more C than the other. Will it, in such situation, always be true that the comparative relation between X and Y changes if one slightly improves either option?

Consider an illustration. Assume, for the sake of the argument, that we can measure health-related quality of life on a scale 0-1. We can illustrate two outcomes that are very different with respect to health-related quality of life and longevity in the following way:

Status quo	Outcome X	Outcome Y
1,000 people at health-related quality of life level .7 with life expectancy 15 years	1,000 people at health-related quality of life level 1 with life expectancy 20 years	1,000 people at health-related quality of life level .7 with life expectancy 28 years

Outcome X is significantly better than Outcome Y with respect to health-related quality of life. If X is chosen, 1,000 people will live at perfect health-related quality of life for 20 years (amounting to 20,000 quality-adjusted life years and 20,000 years of life lived). If Y is chosen, 1,000 people will live at imperfect health-related quality of life but they will live much longer, for 28 years (amounting to 19,600 quality-adjusted life years and 28,000 years of life lived). Insofar as one accepts that neither health-related quality

of life nor longevity is categorically more important than the other, there will be two outcomes like these for which it is true that neither is worse than the other. It might be the one I have presented above, or it might be a different one.

Assuming that the above example is such that neither X nor Y is worse than the other, we can ask the question: Would the comparative relation between the options change if we made a slight improvement of either option? Would Outcome X+ be better than Outcome Y in the table below?

Status quo	Outcome X+	Outcome Y
1,000 people at health-related quality of life level .7 with life expectancy 15 years	1,000 people at health-related quality of life level 1 with life expectancy 20.01 years	1,000 people at health-related quality of life level .7 with life expectancy 28 years

Clearly, Outcome X+ is better than Outcome X. After all, all affected individuals have higher life expectancy in Outcome X+ than in Outcome X. But is it better than Outcome Y? This seems doubtful.

Similarly, would Outcome Y+ be better than Outcome X in the table below?

Status quo	Outcome X	Outcome Y+
1,000 people at health-related quality of life level .7 with life expectancy 15 years	1,000 people at health-related quality of life level 1 with life expectancy 20 years	1,000 people at health-related quality of life level .701 with life expectancy 28 years

Clearly, Outcome Y+ is better than Outcome Y. After all, all affected individuals have higher health-related quality of life in Outcome Y+ than in Outcome Y. But is it better than Outcome X? This seems doubtful.

If it is true that a small improvement of either alternative in this situation fails to change the comparative relation between the alternatives, the covering value (the summary measure of health or its value) is incomplete*. It is incomplete* either because it is vague so that it is indeterminate how some alternatives relate to each other, or because it admits of non-conventional comparative relations such as parity.

As a sidenote, it can be recognized that a similar argument can be made based on the pattern used in the covering value that the cost-effectiveness analysis relies on. Insofar as one, as I myself would endorse, accepts that the pattern must be multidimensional and rely on a combination of patterns proposed by different distributive theories, one can apply the small improvement argument to situations where one option is good in one of the dimensions and another option is good in another dimension. For instance, one can read p as sum total of benefits and r as inequality for an argument targeting the pattern-aspect of the covering value.

The third argument in favor of incompleteness* that I want to present in this paper is that incompleteness* can help one avoid disability-discrimination in certain contexts. Accepting incompleteness* is a way of explaining how certain benefits can be irrelevant.[22] A well-known difficulty in population-level bioethics is that summary measures of health such as QALY tend to imply that it is more valuable to save the life of someone who lives without disabilities than to save the life of someone who lives with disabilities.[23] The reason is straightforward. Saving the life of someone who will live on for 20 years without any disabilities amounts to 20 QALYs, while saving the life of someone who will live on for 20 years with a disability that reduces their health-related quality of life with 0.01 amounts to 19.8 QALYs (20 X 0.99 = 19.8). A QALY maximizer would choose to save the individual living without disabilities. If, instead of relying on QALY maximization as a covering value, we accepted a covering value which was incomplete*, we would be able to say that it is not better to save the life of someone who will live on without disabilities than to save the life of someone who will live with a disability. An incomplete* covering value could, for instance, establish that two such options are on a par.[24]

We have, thus, several good reasons to accept that plausible covering values that can be used together with cost-effectiveness analysis are incomplete*. Plausible covering values are multidimensional, and the

dimensions are so different so that it intuitively seems implausible that the covering values will always manage to fully determine how two items that have some value with respect to it relate to each other. The small improvement argument can be applied to the relevant values that covering values used in cost-effectiveness analysis rely on. Additionally, accepting incompleteness* provides a way of avoiding disability discrimination.

However, it is important to note that incompleteness* in itself does not necessarily pose significant challenges for decision makers.[25] Incompleteness* means that there are some options in some choice sets that cannot be determinately ranked with conventional comparative relations. That is perfectly compatible, at a theoretical level, with there always existing an option that can be determined to be at least as good as all alternatives, and thereby optimal.

When a covering value is incomplete*, it is, however, possible that there will be no available option that is at least as good as all alternatives. This means that it will sometimes be impossible to choose an optimal option (because no such option exists). In extension, this means that one needs to revise decision rules and what criterion one relies on in order to make a choice. Instead of relying on optimization, the following decision rule can be adopted:[26]

> *Determinate maximality*: An option, *x*, is *determinately maximal* with respect to covering value *C* if and only if (a) *x* is not worse than any alternative with respect to *C*, and (b) it is not the case that on all admissible precisifications of *C*, there is an alternative that is better than *x*.

If a covering value admits of non-conventional comparative relations, the first clause, (a), makes sure that options that are not worse than any alternative are considered determinately maximal. This ensures that two options that are, for instance, incomparable with respect to the covering value are both determinately maximal although they are not optimal (at least as good as all options) as long as there are no alternatives that are better than them in the choice set. This clause is inspired by how Amartya Sen claims that a rational choice is a choice that is maximal.[27] If a covering value is vague and fails to determine which of the conventional comparative relations (better than, worse than, equally as good as) obtains between all options in all situations, clause (b) ensures that options for which it can be determined that they are worse than some alternative are not considered determinately maximal. This condition is inspired by how those who believe incommensurability is vagueness deals with decision problems by invoking supervaluationism.[28] According to the supervaluationist approach to vagueness, a statement is supertrue (superfalse) if and only if it is true (false) on all admissible precisifications, and an admissible precisification is a precisification that specifies a vague term in accordance with normal language and respects penumbral truths.[29] For instance, it is admissible to precisify the term bald so that the king of Sweden is bald, but not admissible to precisify the term bald so that the queen of Sweden is bald. Together, the two clauses allow one to express a decision criterion that can be used together with a covering value that is incomplete* without taking a stand on what the explanation of the incompleteness* is (non-conventional comparative relations or vagueness or both).

It is obvious that all optimal options, that is, options that are at least as good as every alternative, are also determinately maximal options. However, when covering values are incomplete*, not all determinately maximal options will be optimal.

Relying only on determinate maximality as a decision criterion when covering values are incomplete* has some theoretical problems.[30] For instance, it can lead to so-called dynamic choice inconsistencies and a kind of money pump. To see this, consider someone choosing first between X, X+ and Y, and after that gets an opportunity to change the decision – whatever it is – to X. If determinate maximality is the only decision criterion, it would be rational to first choose Y and then change one's mind to X (Y is determinately maximal in the set {X, X+, Y} and X is determinately maximal in the set {X, Y}. However, X is determinately worse than some available option that the agent could have chosen, namely X+. Nevertheless, determinate maximality is a way of avoiding being paralyzed by incompleteness*. As such, it is an important tool that can be used to help decision makers choose among options when some alternatives cannot be ranked.

Discrimination Risk

In this section, I will outline how incompleteness* introduces an increased risk of discrimination compared to situations in which appropriate covering values always manage to fully determine at least one option that is at least as good as all alternatives.

Since appropriate covering values in the healthcare sector will be incomplete*, those using such covering values in cost-effectiveness analysis will find that the decision tool is sometimes unable to determine an optimal option. Sometimes, there will be no option that can be determined to be at least as good as all alternatives. As explained in the previous section, this does not paralyze decision makers since decision makers can replace optimality requirements with something like determinate maximality. Nevertheless, it has several implications.

First, it means that even if cost-effectiveness analysis combined with determinate maximality as a criterion for choice can be used to always identify a set of justified options, this set of justified options will often be larger than if cost-effectiveness analysis had been able to always fully determine at least one option that is at least as good as all alternatives. On the one hand, there will be more situations in which more than one option is determinately maximal compared to the number of situations that are characterized by the fact that there is more than one option that is optimal. On the other hand, the set of options that are determinately maximal will often be larger than it would have been if cost-effectiveness analysis had been able to always fully determine an option that is optimal.

Second, the incompleteness* entails that the justificatory power of the cost-effectiveness analysis is reduced.[31] If cost-effectiveness analysis always helped decision makers identify an optimal option, cost-effectiveness analysis would always provide a very strong justification for choosing an option; there would be at least one option that is best. That is a strong and often sufficient justification for choosing one of those options. If cost-effectiveness analysis sometimes is only able to identify alternatives that are determinately maximal but not optimal, the justification for choosing an option would sometimes be significantly reduced. Ruth Chang has suggested that what happens when one faces two options that are determinately maximal while not being optimal is that a "resolutional remainder" arises.[32] As the covering value one uses for the comparison fails to fully determine an optimal option, it leaves the decision maker with something unresolved. If cost-effectiveness analysis actualizes resolutional remainder of this kind, the justification for choosing a determinately maximal option is sometimes insufficiently strong.

The implications of the incompleteness* is an increased risk of discrimination. The reason is straightforward. When several options are deemed permissible (either because they are optimal or because they are determinately maximal), something other than cost-effectiveness analysis will determine what the decision maker eventually chooses. This can be many things, for instance, self-interest, biases, whims, arbitrariness, or the decision maker's own assessments. Importantly, whatever it is that in the end determines what the decision maker chooses, there is a risk that the choices that are made in these situations will systematically benefit some groups over others; there is a risk of discrimination.[33]

Dealing with Discrimination Risk

The discrimination risk that follows from incomplete* appropriate covering values in cost-effectiveness analysis in the healthcare sector can be dealt with in different ways. I will discuss the two most obvious kinds of approaches: (i) to hold on to a cost-effectiveness analysis that always determines that there is at least one option that is optimal and embrace a covering value that has shortcomings, and (ii) to accept that cost-effectiveness analysis will be incomplete* because it must rely on a covering value that is incomplete*. The first approach means that one embraces and uses some covering value that one knows gives certain health conditions and certain groups in the population some unwarranted relative advantage over other health conditions and other groups in the population when it comes to chances of receiving resources by the allocator. The second approach leads to some technical problems for those designing cost-effectiveness analyses and also entails that the decision tool leaves more options to the decision maker. Both approaches actualize discrimination problems.

Let us start with the first approach: using a covering value that is not incomplete* although one has good reason to believe that all appropriate covering values are incomplete*. In effect, this is what happens when health economists and others rely on QALY, DALY, and other summary measures of health together with cost-effectiveness analysis. These summary measures of health are not vague, and they do not admit of incomparability, incommensurability or parity. Instead, they always fully determine how all options relate to each other. If, however, I am right that we ought to expect that all appropriate covering values that can be used for valuations of health states and health interventions are incomplete*, we know that all of these summary measures of health and summary measures of the value of health are flawed. They have formal properties that we know an appropriate covering value cannot have. This means that we know that these summary measures systematically overvalue certain health states and health interventions, and we know that they systematically undervalue other health states and health interventions. We might not know which health states and health interventions are overvalued and which are undervalued, but we know that some are overvalued and that some are undervalued. This follows directly from the fact that they always fully determine a ranking of all health states and all health interventions under conditions where we know that an appropriate covering value is a covering value that cannot fully determine a ranking of all health states and all health interventions. Essentially, this approach means that one places the discrimination risk in the choice of covering value.

The second approach leads to other problems. First, it requires health economists and policy makers to think differently about cost-effectiveness analysis and comparisons. Instead of expecting cost-effectiveness and comparisons to generate conventional rankings, health economists and policy makers must, according to this approach, adopt a way of thinking according to which the purpose of cost-effectiveness analysis and comparisons more broadly is to identify ineligible options. Second, as outlined in the previous section, this approach entails that decision makers who rely on cost-effectiveness analysis will often have to choose between many more options after the cost-effectiveness analysis has eliminated the ineligible options. Essentially, this approach means that one places the discrimination risk in the hands of the decision maker who uses cost-effectiveness analysis.

Corresponding to the two types of approaches one might take to the incompleteness*, there are two types of approaches to the discrimination risk. First, one can address the discrimination risk by embracing a covering value that always fully determines an option that is at least as good as all alternatives and attempt to reduce the amount of discrimination that might follow by tinkering with the covering value. If, for instance, one accepts that QALY by necessity (due to incompleteness*) will get some valuations wrong, one can try to balance the implications of these wrongs so that one avoids disproportionate distributions of benefits between socially salient groups by changing the valuations of health states that are prevalent in different socially salient groups. For example, if back pain is more prevalent in men than in women and if one has reason to believe that men disproportionally benefit from the healthcare system, one can adjust the disvalue associated with back pain. Similar moves can, in principle, be made for all health conditions and all socially salient groups.

Second, one can address the discrimination risk by embracing an incomplete* covering value and set of guidelines for decision makers with the purpose of reducing discrimination. Such anti-discrimination guidelines can be designed in different ways. One type of anti-discrimination guideline would focus on equalizing the distribution of benefits across salient social groups. A different type of anti-discrimination guideline instead focuses on making certain grounds impermissible to use when choosing between determinately maximal options. Neither of these types of guidelines will be easy to define in detail and implement, but they are at least possible pathways forward.

Regardless of how one approaches discrimination risks that arise due to the fact that appropriate covering values in the healthcare sector are incomplete*, it will be difficult to completely safeguard against discrimination. However, by recognizing the discrimination risks that exist, I hope that this paper can help raise awareness of these problems*.

Acknowledgments. I am grateful for the comments from an anonymous reviewer. Financial support was provided from the Swedish Research Council, grant number 2017–01382.

Notes

1. Ord T. The moral imperative toward cost-effectiveness in global health. *Center for Global Development* 2013;12; Hausman D. *Valuing Health: Well-Being, Freedom, and Suffering*. New York: Oxford University Press; 2015.
2. Bognar G, Hirose I. *The Ethics of Health Care Rationing.* New York: Routledge; 2014; Cookson R. Justice and the NICE approach. *Journal of Medical Ethics* 2015;**41**:99–102. See also note 1, Ord 2013 and Hausman 2015.
3. Dolan P. The measurement of individual utility and social welfare. *Journal of Health Economics* 1998;17:39–52; Bleichrodt H, Diecidue E, Quiggin J. Equity weights in the allocation of health care: the rank-dependent QALY model. *Journal of Health Economics* 2004;**23**:157–171. Nord E. Concerns for the worse off: Fair innings versus severity. *Social Science and Medicine* 2005;**60**:257–263; Bognar G. Age-weighting. Economics and Philosophy 2008;24:167–189; Herlitz A, Horan D. Measuring needs for priority setting in healthcare planning and policy. *Social Science and Medicine* 2016;157:96–102; Herlitz A. Income-based equity weights in health care planning and policy. *Journal of Medical Ethics* 2017;43:510–514.
4. Chang R. Introduction. In: Chang R, ed. *Incommensurability, Incomparability, and Practical Reason.* Cambridge: Harvard University Press; 1997; Chang R. The possibility of parity. *Ethics* 2002;112:659-688; Andersson H, Herlitz A. Classifying comparability problems in a way that matters. The Institute for Futures Studies. *Working* Paper 2021:22; Andersson H, Herlitz A. Introduction. In: Andersson H, Herlitz A, eds. *Value Incommensurability: Ethics, Risk, and Decision-Making.* New York: Routledge; 2022.
5. Eyal N, Herlitz A. Symposium: Ethics of economic ordeals. *Economics & Philosophy* 2021;**37**:1–7.
6. Sen A. 1980. Equality of what? In: McMurrin S, ed. *Tanner lectures on human values*. Cambridge: Cambridge University Press; 1980; Cohen GA. On the currency of egalitarian justice. *Ethics* 1989;99:906–944.
7. Nozick R. *Anarchy, State and Utopia.* New York: Basic Books; 1974.
8. For a good discussion of these, and an argument for a summary measure of the value of health, see note 1, Hausman 2015.
9. Tännsjö T. *Setting Health-Care Priorities: What Ethical Theories Tell Us*. Oxford: Oxford University Press; 2018. See also note 1, Ord 2013.
10. Ottersen T. Lifetime QALY prioritarianism in priority setting. *Journal of Medical Ethics* 2013;**39**:175–180; Eyal N, Herlitz A.. Input and output in distributive ethics. *Noûs* DOI:10.1111/nous.12392
11. Harris J. *The Value of Life.* London: Routledge & Kegan Paul; 1985; Herlitz A. The indispensability of sufficientarianism. *Critical Review of International Social and Political Philosophy* 2019;**22**:929–942.
12. Temkin L. Egalitarianism defended. *Ethics* 2003;113:764–782; Eyal N, Hurst SA, Norheim OF, Wikler D, eds. *Inequalities in Health: Concepts, Measures, and Ethics*. Oxford: Oxford University Press; 2013.
13. See note 4, Andersson, Herlitz 2022.
14. Broome J. *Is incommensurability vagueness?* In: Chang R, ed. *Incommensurability, Incomparability, and Practical Reason*. Cambridge: Harvard University Press; 1997; Broome J. Incommensurability is vagueness. In: Andersson H, Herlitz A, eds. *Value Incommensurability: Ethics, Risk, and Decision-Making.* New York: Routledge; 2022.
15. Chang, R. Are hard cases vague cases? In: Andersson H, Herlitz A, eds. *Value Incommensurability: Ethics, Risk, and Decision-Making*. New York: Routledge; 2022; Rabinowicz W. Value relations. *Theoria* 2008;74: 18-49. See also note 4, Chang 2002.
16. See note 4, Chang 2002, Andersson, Herlitz 2021, 2022, note 14, Broome 1997, 2022, note 15, Rabinowicz 2008, Chang 2022.
17. On non-comparability, see note 4, Andersson, Herlitz 2022.
18. Savage LJ. *The Foundations of Statistics* Hoboken: Wiley; 1954. See also note 4, Andersson, Herlitz 2022.

19. Raz J. *The Morality of Freedom* Oxford: Clarendon Press; 1986; Qizilbash M. The mere addition paradox, parity and vagueness. *Philosophy and Phenomenological Research* 2007;**75**:129–151. See also, note 4, Chang 2002, Andersson, Herlitz 2021, 2022.
20. See note 4, Chang 2002.
21. Herlitz A. Indeterminacy and the principle of need. *Theoretical Medicine and Bioethics* 2017;**38**:1–14.
22. Kamm F. Aggregation, allocating scarce resources, and the disabled. *Social Philosophy and Policy* 2009;**26**:148–197.
23. Beckstead N Ord T. Bubbles under the wallpaper: Healthcare rationing and discrimination. In Kuhse H, Schüklenk U, Singer P, eds. *Bioethics: An Anthology.* Oxford: Blackwell. 2015;406–12; Hassoun N, Herlitz A. Distributing global health resources: Contemporary issues in political philosophy. *Philosophy Compass* 2019;14.
24. See note 22, Kamm 2009.
25. Herlitz A. The limited impact of indeterminacy for healthcare rationing: How indeterminacy problems show the need for a hybrid theory, but nothing more. *Journal of Medical Ethics* 2016;**42**:22–25.
26. Herlitz A. Nondeterminacy, cycles and rational choice. *Analysis* 2020;**80**:443–449; Herlitz A, Sadek K. Social choice, nondeterminacy, and public reasoning. Res Philosophica 2021;98:1–25; Herlitz A. Nondeterminacy and reasonable choice. In Andersson H, Herlitz A, eds. *Value Incommensurability: Ethics, Risk, and Decision-Making.* New York: Routledge; 2022; See also note 4, Andersson, Herlitz 2021.
27. Sen A. Maximization and the act of choice. *Econometrica* 1997;**65**:745–779.
28. Broome J. Reply to Rabinowicz. Philosophical Issues 2009;19:412–417; Herlitz A. Nondeterminacy, two-step models, and justified choice. *Ethics* 2019;**129**:284–308.
29. Fine K. Vagueness, truth and logic. *Synthese* 1975;**30**:265–300.
30. See note 4, Andersson, Herlitz 2022.
31. See note 26, Herlitz 2022.
32. See note 4, Chang 2002, Andersson, Herlitz 2021, note 15, Chang 2022. See also Wasserman R. Indeterminacy, ignorance and the possibility of parity. *Philosophical Perspectives* 2004;**18**:391–403.
33. Herlitz A. Microlevel prioritization and incommensurability. *Cambridge Quarterly of Healthcare Ethics* 2018;**27**:75–86.

Cite this article: Herlitz A (2023). Cost-Effectiveness, Incompleteness, and Discrimination. *Cambridge Quarterly of Healthcare Ethics* 32: 163–173, doi:10.1017/S0963180122000263