This paper has been published as: Walker M and Rogers W. Defining disease in the context of overdiagnosis. *Medicine, Healthcare and Philosophy* 2017; 20 (2): 269-280. DOI: 10.1007/s11019-016-9748-8

Defining disease in the context of overdiagnosis

Mary Jean Walker and Wendy Rogers

Recently, concerns have been raised about the phenomenon of 'overdiagnosis'. This is, roughly, the diagnosis of a condition that is not causing harm, and will not come to cause harm.[1] It differs from misdiagnosis (including false positives) in that it does not involve an identifiable mistake, as current diagnostic criteria are met. Overdiagnosis is problematic because patients receive diagnoses that do not benefit, or may harm them, by the imposition of medical labels and/or resulting unnecessary interventions. Concern about overdiagnosis is linked to concerns about overtreatment, medicalisation, and disease mongering, but the term refers specifically to the problem of diagnosing non-harmful disease.[2] Overdiagnosis occurs in relation to a range of common diseases, including breast, thyroid, and prostate cancers, type 2 diabetes, and osteoporosis, among others.[3]

Overdiagnosis raises a host of questions, including ethical questions: how do we weigh the harms to some patients from overdiagnosis against the benefits to others from the same preventative strategies?; practical questions: ought we reduce screening programs?; and scientific questions: are there ways to tell which cases of a disease are likely to progress to cause harm, and which are not? It also raises questions about disease definition.

[1] How to define overdiagnosis is under dispute. Welch et al (2011, xiv) define it as occurring "when individuals are diagnosed with conditions that will never cause symptoms or death". Carter et al define it by saying "An (asymptomatic) person is diagnosed with a condition; that diagnosis does not produce a net benefit for that person" (2015). The reference to net benefit in some definitions may be problematic since 'diagnoses that do not result in net benefit' would capture a broader range of cases where there are unintended iatrogenic harms (see Rogers and Mintzker 2016). Some also allow that overdiagnosis can happen to patients with symptoms. We clarify this rough statement in the discussion in section 1.

[2] The term has sometimes been used in a broad sense, to include a set of related problems such as medicalisation, and overtreatment, as well as in the narrower sense we focus on here. See Carter et al (2015); Moynihan, Doust, and Henry (2012).

[3] We limit our discussion to physical disease, as psychiatric disease raises further complexities that are outside our current scope.

This paper has been published as: Walker M and Rogers W. Defining disease in the context of overdiagnosis. *Medicine, Healthcare and Philosophy* 2017; 20 (2): 269-280. DOI: 10.1007/s11019-016-9748-8

Overdiagnosis challenges how we define specific diseases, since where a disease is overdiagnosed, this is evidence that current diagnostic criteria (the clinical or medical definition of those particular diseases, and what we take as evidence that the criteria are met) are inadequate. But overdiagnosis also raises questions about defining our concept 'disease', and these will be our focus in this paper. We show how overdiagnosis challenges existing accounts of the concept 'disease', argue that we should regard 'disease' as a vague concept, and then seek to develop a précising definition of disease that can help to prevent and limit overdiagnosis.

In section 1, we explain overdiagnosis, and show how it raises questions about the way that disease is defined. In section 2 we examine how several influential accounts of disease might accommodate the phenomenon of overdiagnosis. None seem able to, in part at least because of a shared underlying assumption that disease is a classically structured concept with clear boundaries. We argue, in section 3, that disease should instead be considered a vague concept. On this view, overdiagnosed cases are 'borderline' cases of disease, falling in the zone between cases that are clearly disease, and cases that are clearly not disease. While overdiagnosed cases are therefore not clearly categorised by our concept 'disease', they can be sorted by using précising definitions of disease.

In section 4, we develop a précising definition of disease, designed to be of use in preventing and limiting overdiagnosis. We argue that for this purpose, we can define disease as physiological dysfunction that has a significant risk of causing severe harm to the patient. These criteria ('significant risk' and 'severe harm') will require specification in different cases, taking account of the best current information available on the condition, societal views, and individual patient preferences. This approach results in a definition that is relative

to patient's views about what counts as a severe harm, and their preferences regarding the level of risk they are willing to accept, exercised within parameters that limit what may be counted as disease.

1. Overdiagnosis

Population statistics indicate that for some diseases, finding (and treating) an increased number of cases does not reduce mortality and/or morbidity associated with that disease. In recent years increased implementation and reach of screening programs, and the introduction of more sensitive testing and imaging technologies, have triggered increases in the rates of identified disease. For example, breast cancer incidence increased by an average of 2-10% per year among 50-69 year-old women after the implementation of mammography screening programs in the USA, UK, Europe, and Australia (Barratt 2015). The increase primarily reflects the fact that screening identifies more cases that are less serious or advanced – cases that were not detectable without sophisticated imaging. Cases that are clinically evident, or causing symptoms, are of course detectable without such methods.

The rationale for screening is that finding less advanced cases will enable interventions to prevent more severe, advanced disease. As such we would expect that finding more of these less serious cases would lead to falls in morbidity and mortality from these diseases, as treatment is often more effective earlier in the course of disease. We would also expect (at least after screening has been in place for some years) to find fewer cases of more advanced disease, as cases will have been identified and treated earlier in their development. But neither of these results has occurred: both mortality rates and incidence of advanced cases

are, for some diseases and populations, not substantially affected by screening programs.[4] We see a similar pattern of increasing rates of diagnosis but no concomitant fall in mortality in conditions diagnosed using new, highly sensitive testing modalities, outside of screening programs. For instance, after the introduction of high resolution CT pulmonary angiography for detecting pulmonary emboli, there was an 80% increase in incidence, with no significant change in mortality from PE, over the 8 years following its introduction in the USA (Hoda 2014).

One explanation for such statistical stories is that we are finding more cases because there have been actual spikes in the rates of these diseases. But if this were the case, we would expect mortality to increase a corresponding amount (unless we posit that treatment happens to have improved simultaneously, and to an exact corresponding extent), which has not occurred (Welch et al 2011, 55-6). A more plausible explanation is that increasing the detection and treatment of many of these conditions has no effect on outcomes because the additional cases found would not have progressed to more serious disease or caused patients harm, had they remained undetected. This reflects the notion that there is a 'reservoir' of cases of sub-clinical disease, that either never progress to become clinically evident, or are self-limiting, and which remained undetected prior to the advent of screening and more sophisticated diagnostic testing. These are the cases that, when found and diagnosed alongside progressive cases of disease, are overdiagnoses.

This explanation of the statistical stories receives confirmation from autopsy studies, and

[4] For example, while there is some evidence indicating that mammography screening programs can reduce deaths from breast cancer when targeted at women aged 50-69 (Barratt 2015), this is not the case for breast cancer screening that involves women younger that 50 (Autier 2015) or older than 70 (Glas et al 2014).

screening studies of healthy people. Autopsy studies of people who died from unrelated causes have found that 1.3% of women had previously unsuspected breast cancer, and 9% had ductal carcinoma in situ (DCIS), a non-invasive lesion that may or may not progress to an invasive stage (Barratt 2015). In relation to prostate cancer: "nearly 10% of men in their 20s and 80% of men in their 70s have prostate cancer discovered incidentally on autopsy after they die from an accident, yet only 3% of men die of prostate cancer" (Coon et al 2014, 2). The population prevalence of microscopic papillary thyroid cancers is thought, based on autopsy findings, to approach 100 per cent, although less than 0.1% of these will be clinically apparent (Harach et al 1985). Similarly, studies that have performed CT scans of symptom-free, healthy people have concluded that 'the average individual had 2.8 abnormalities' (Welch et al 2011, 38).

Overdiagnosis also occurs in relation to 'spectral' diseases, which are defined using a threshold point on measures of biological variables, such as type 2 diabetes, osteoporosis, chronic kidney disease, and hypertension. With these conditions, increases in incidence are at least partly due to intentional changes to the diagnostic criteria which shift the threshold point distinguishing a 'normal' result to one indicative of disease. For example, the diagnostic threshold for diabetes, previously 140 mg/dL fasting blood glucose, was defined at 126 mg/dL fasting blood glucose in 1997; the threshold for diagnosing hypercholesterolaemia was altered from 240 to 200 mg/dL of blood cholesterol in 1998; the threshold for diagnosing hypertension was altered from 160 mm Hg systolic or 100 mm Hg diastolic, to 140 and 90 respectively, in 1997 (Moynihan et al 2013).

The reasoning behind these changes is, in part, similar to the motivation for screening: a more expansive definition will enable identification of disease when it is less severe, and so

facilitate interventions to prevent morbidity associated with more advanced disease. The diagnostic thresholds are set in ways that reflect information about the risks these conditions pose, and the potential reduction of risk through treatment. But as the thresholds lower, the chance that a particular person will benefit from treatment decreases, and the number of people that need to be treated to benefit one person increases (Welch et al 2011, ch 1). For instance, studies comparing outcomes for people with hypertension with and without treatment provide the following statistics. For those with a diastolic blood pressure over 115 (severe hypertension), 72% of those treated will benefit (i.e., will avoid a serious complication of hypertension, such as stroke or heart attack). This means that 1.4 people need to be treated for one person to benefit. In contrast, for those with a diastolic blood pressure of 90-100 (mild hypertension), only 6% of those treated will benefit. Eighteen people need to be treated in order for one person to benefit (Welch et al 2011, 6). Many of those who have the disease would not have been harmed by it in the absence of diagnosis or treatment; that is, with lower thresholds, more people are overdiagnosed.

Rogers and Mintzker (2016) propose a taxonomy of overdiagnosis that distinguishes between "maldetection" and "misclassification", which maps loosely on to the two explanations of the phenomenon given above. Maldetection overdiagnoses occur where we are unable to discriminate between apparently identical lesions or abnormalities, only some of which will lead to harm, as in the case of screening-detected cancer overdiagnoses. Misclassification overdiagnoses arise in relation to the lowering of diagnostic thresholds for diseases defined in terms of physiological parameters that occur on a spectrum, such that many with a positive diagnosis do not experience symptomatic disease.[5]

[5] These are not mutually exclusive types as there are cases where both misclassification and maldetection occur, such as pulmonary embolism.

This paper has been published as: Walker M and Rogers W. Defining disease in the context of overdiagnosis. *Medicine, Healthcare and Philosophy* 2017; 20 (2): 269-280. DOI: 10.1007/s11019-016-9748-8

Overdiagnosis can result in harms for patients. It leads to unnecessary interventions, some of which are highly invasive and/or have significant side effects. It causes unnecessary costs for healthcare systems, and may contribute to inequalities in healthcare resource distribution (Godlee 2012; Moynihan, Henry and Moons 2014). The problem of overdiagnosis does not reduce to these problems of overtreatment (i.e., unnecessary interventions); diagnosing someone – giving them a medical label and the status of a patient – itself has effects that may cause suffering ([redacted]). There are negative psychological effects of being labelled with a disease (Coon et al 2014; Pickering 2006); some have argued that there are negative cultural effects related to increased levels of medical anxiety and surveillance (Armstrong 1995). But in addition to these detrimental practical consequences, overdiagnosis also involves epistemic problems.

Overdiagnosis arises, in part, due to epistemic gaps – we do not have enough knowledge. We detect an abnormality – a lesion, or an abnormal reading of some biological variable – but do not know how to distinguish cases where the abnormality will progress to harmful disease from those where it will not.[6] A characteristic of overdiagnosis is that although we know at a population level that overdiagnosis occurs, the information we need to determine contemporaneously which individuals are overdiagnosed will only be available in the future, or counterfactually (Hofmann 2013). The only way to be certain that a particular person is overdiagnosed would be to ignore the diagnosis, and wait until after their death to see whether or not they were harmed by the condition (Welch et al 2011, 55). To fill this

[6] This is perhaps the difference between overdiagnosis and other cases where a diagnosis causes overall net harm to a patient (see note 1): in overdiagnosis, the condition diagnosed seems to be the same as cases where a diagnosis will benefit the patient. This could indicate that it is the epistemic problems that are definitive of overdiagnosis, though we do not argue this here.

epistemic gap we could increase our knowledge. For instance, genomics may be able to provide some ways of predicting which cancers will progress to become harmful, and perhaps future diagnostic tests will be able to distinguish which cases will progress, for other diseases. But in the absence of any capacity to distinguish which of the currently diagnosed cases are overdiagnoses, it seems that overdiagnosis involves epistemic uncertainty.

But overdiagnosis could not be completely remedied by increasing our knowledge in these ways, because it involves not only epistemic uncertainty, but metaphysical indeterminacy.[7] There may be in principle no way to distinguish, at the time of diagnosis, whether a diagnosis is an overdiagnosis, because this may be metaphysically indeterminate. Whether a particular case of disease progresses or not may be a matter of chance (e.g., chance mutations acquired by clonal cells that lead a group of abnormal cells to develop on a malignant trajectory); or may depend on what we decide to do now (including whether or not there is intervention). The epistemic uncertainty reflects not just lack of knowledge, but that the way a disease progresses or not, can be metaphysically indeterminate.

Thus, there is a sense in which overdiagnosis involves having too much information, or more precisely, information that we don't know (even though we thought we knew) how to classify or respond to. The results of some diagnostic tests seem mere 'noise' rather than information, since they seem to provide no useful direction (Welch et al 2011, ch 13; Rogers and Mintzker

[7] Here we are using 'metaphysical indeterminacy' to refer to open causal chains from the currently identified abnormal state, and its future progression or non-progression (we ignore the possibility of metaphysical determinism for the purposes of this paper). We note that this differs from the way that Hofmann (2016) uses 'indeterminacy' to refer to the social question of whether or not a condition comes to be considered a disease through the process of medicalisation (i.e. whether some conditions are diseases is resolved by social and political processes rather than by appeal to underlying biological processes). Thus we are not here disputing Hofmann's analysis of the differences between medicalisation and overdiagnosis.

2016, 585).[8] Further, the evidence for overdiagnosis disrupts what we thought we knew about which dysfunctions or abnormalities are linked to what symptoms and effects. That is, overdiagnosis raises questions about how we classify states into healthy and pathological, normal and abnormal.

2. How do existing accounts of disease accommodate overdiagnosis?

The discussion in section 1 demonstrates that overdiagnosis raises questions about a number of features commonly linked to the concept 'disease'. Specifically, overdiagnosis involves the presence of physiological dysfunction[9] that meets the criteria for identifying specific diseases, but without the harmful sequelae usually or often associated with that disease. Thus it raises questions about the conceptual links between disease and dysfunction, harm and risk.[10] To investigate the challenge of overdiagnosis for our notion of disease, in this section we examine how a number of influential accounts of disease based on dysfunction and/or harm might categorise overdiagnosed cases.[11] We focus on Boorse's biostatistical theory (BST); Nordenfelt's view of disease as an incapacity to realise one's vital goals; Wakefield's Harmful Dysfunction Analysis (HDA); and Schwartz's discussion of risk and disease. Our aim in this section is to clarify how overdiagnosis challenges conceptual links between disease and dysfunction, harm, and increased risk, via examination of these theories.

The BST defines disease as an internal state that reduces functioning to a level below that which is statistically typical for the species (Boorse 1977, 555). The function of a cell, organ,

[8] We thank an anonymous reviewer for reminding us of this point.
[9] We use 'dysfunction' in Boorse's naturalistic sense here and throughout the paper.
[10] It also challenges the link to the imperative to intervene that some have emphasised (e.g., Engelhardt 1975). We leave this aside in this section, as we consider it linked to the features of harm and increased risk of harm.
[11] We will ultimately depart from these theories, in arguing in section 3 that 'disease' is a vague concept that is not classically structured.

or organ system is its species-typical contribution to some biological goal (e.g., digestion; circulation). The notion of 'dysfunction' has become the standard way to capture the sense that disease involves something 'going wrong' with the body, a departure from what is normal or 'meant' to happen, in a naturalistic sense. The corresponding notion of 'function' can be fleshed out either aetiologically (in terms of evolutionary theory) or, as Boorse does, in terms of contributions to biological goals (1976).

On the BST, all cases of maldetection overdiagnosis (where an abnormality is detected that may be either aggressive or indolent) are disease. Any abnormal lesion is a dysfunction since it involves cells, organs, or organ systems failing to make their species-typical contribution to biological goals. The inclusivity of the BST in this regard reflects Boorse's claim that harm is not conceptually connected to disease. His reasoning is that cases of disease that are not harmful, such as asymptomatic, or very mild self-limiting diseases, show that harm is not necessary for disease (Boorse 1977, 547).[12]

Regarding misclassification overdiagnosis (where disease thresholds are lowered to include ever smaller deviations from normal function), the BST does not provide a principled way to decide on the level of dysfunction at which we count some condition as disease. Although it appeals to statistical abnormality, the exact level of abnormality required for disease is

[12] This feature of the BST is demonstrated in the 'one dead cell' objection: that the BST implies disease is present where there is just one dead cell, since that cell is failing to make its species-typical contribution to the biological goals of the organism, even if there is no noticeable change in overall functioning (Nordenfelt 1987, 28, cited in Wakefield 2014, 656). Boorse accepts this implication for tissues that do not usually involve regular cell death and regeneration (1987, 365, in Wakefield 2014, 656). For tissues that do involve regular cell death and regeneration, he indicates that one dead cell would not be pathological since it is not statistically atypical (1997, 50-51). Correlatively, one might say that if it turns out to be normal to have microscopic thyroid cancers, or other lesions, they are not disease on the BST. However, Boorse also accepts that there are some 'universal' diseases, such as tooth decay, and his ways of dealing with those would also apply to universal maldetected cancers. We leave further discussion of this aside since our focus is assessing the relevance of dysfunction to disease, not statistical frequency.

necessarily somewhat arbitrary and "can only be conventionally chosen" (1977, 559). Boorse states that this problem is merely "academic" because "most diseases involve functional deficits that are unusual by any reasonable standard" (1977, 559).[13] Thus with regard to spectral diseases, the BST faces a threshold problem, and cannot specify in any principled way where on the spectrum 'dysfunction' begins (see Schwartz 2007a). Part of the difficulty here is that whatever threshold is set, there will be some people above it who receive a diagnosis that does not benefit them, and some people below it who would have benefitted from a diagnosis. The threshold problem reflects the challenge of categorising cases as disease or non-disease (i.e., a binary distinction), given that the relevant 'dysfunction' can occur by degrees ([redacted]). Thus the BST accepts maldetection overdiagnoses as disease, and cannot adjudicate on the question of misclassification overdiagnoses.

Other approaches to defining disease begin not from dysfunction, but from the effect that disease has on people's lives, or what we can loosely call harm.[14] If we are dissatisfied with the BST's accommodation of overdiagnosis, harm would seem to be the relevant concept to which to appeal. Nordenfelt proposes that health is the ability to realise one's 'vital goals' (that is, goals related to surviving and flourishing), and disease is a bodily or mental process which tends to reduce its bearer's health (1993, 280). Since cases of (both misclassification and maldetection) overdiagnosis do not harm their bearer, it might seem that on this account they are not disease. But for Nordenfelt, "The process need not actually reduce his or her

[13] This is not a satisfying answer since on his theory what counts as a functional deficit is supposed to respond only to the standard of statistical frequency (for more detail on this problem, see Schwartz 2007a). In any case, there does not seem any obvious cut-off point at which the dysfunction involved should be thought unusual enough to indicate the presence of disease.

[14] 'Harm' may of course refer to a range of things, including pain and suffering, incapacity, or presence of symptoms. Though we recognise that there are various ways these things might come apart, we will discuss incapacity, pain and suffering simply as 'harms' in what follows for reasons of space.

health. The disease can be latent or it can have aborted at an early stage. It is, however, still called a disease because it belongs to a type, of which the majority of instances actually reduce the health of their bearers" (1993, 280). This might indicate that overdiagnoses are disease, on Nordenfelt's account. However, the evidence for overdiagnosis makes it unclear whether or not the types to which the cases belong do cause harm in a 'majority' of cases. For many diseases, we do not know the size of the reservoir of non-harmful cases, so that we do not know if the majority of these cases reduce the health of their bearer. Nordenfelt also does not discuss what proportion of the total cases would count as a majority, leading to a threshold problem. Appeal to 'a majority' seems to conflict with Nordenfelt's eschewal, elsewhere, of statistical notions of normality. There may also be issues here in relation to how we define disease 'types': is DCIS itself a 'type', or would it be a subtype of lesions of breast cancer? Finally, harm is a matter of degree, so that accounts appealing to harm will also face threshold issues of deciding the level of harm indicative of disease (a point Nordenfelt acknowledges). Thus other features of his account complicate an initial view that Nordenfelt would classify overdiagnoses as disease.

Wakefield's HDA considers disease[15] necessarily to involve both dysfunction and harm. Wakefield rejects dysfunction as individually sufficient for disease because, he says, having a dysfunction is consistent with being completely healthy (2014, 655). He demonstrates this with examples of individuals with dysfunctions who do not experience harm: asymptomatic disease carriers, and others with neutral, risky, or benign mutations. He shows that such cases are considered (by the medical profession, and established nosology) to be healthy rather than diseased, unless or until their condition involves harms (2014, 668). Wakefield's response to

[15] Wakefield uses 'disorder' rather than 'disease', but states that the term fulfils a similar role in his account to 'disease' in other parts of the debate. We thus alter his terminology here for consistency.

the apparent existence of diseases that do not cause harm – the asymptomatic, or mild self-limiting cases – is that such diseases are benign, or less harmful, versions of states that *usually* are harmful (thus the need to qualify 'disease' with another term such as 'benign' to specify what we mean) (2014, 673-4). That is, benign diseases are not, properly speaking, diseases at all, but rather a separate and parasitic concept, whose existence does not count as a counterexample to his analysis.

On the HDA, it seems that both maldetection and misclassification overdiagnosis would not count as diseases. Overdiagnosed patients are similar to asymptomatic disease carriers, or people with benign mutations: they have a physiological dysfunction, but experience no harm, and thus do not have a disease. However, in the context of overdiagnosis, the requirement for specific cases to involve harm to the bearer makes the HDA overly narrow. Someone with colonic polyp with malignant changes but no local or distant invasion, and no experienced harm, nonetheless does have a disease. Wakefield's response to the notion of non-harmful diseases, that they are not really diseases, seems wrong in relation to this kind of example, even if it is plausible in relation to asymptomatic disease carriers and genetic mutations.[16] The issue here is perhaps that Wakefield aims, we take it, to rule out as disease cases where there is *and will be* no harm, but perhaps he would wish to include as disease cases where there is no current harm, but *future* harm is possible or likely. And it is future harm that is at issue for deciding whether overdiagnosis is disease. Wakefield does not explicitly address the role of future harms in his theory.

[16] One might also think that the presence of such dysfunction itself constitutes a harm. We leave this aside since it could lead to an elision of the two criteria. Thus we are considering harm as an evaluative and not biological notion.

Schwartz (2008) considers future harm as a criterion for disease in his 'function and negative consequences' (FNC) view, which adds a 'negative consequences' requirement to Boorse's statistical approach. This would be consistent with the way that some specific diseases are defined in relation to risks. For instance, the level of fasting blood glucose selected to diagnose type 2 diabetes reflects epidemiological information related to the risk of developing diabetic retinopathy. The level of blood pressure taken to indicate hypertension reflects available risk information on levels of blood pressure associated with heart attack, stroke, and microvascular damage. Schwartz (2008) points out that some risk factors have come to be treated as diseases because the risk of future harm has come to influence what we categorise as disease.

But as Schwartz argues, using risk of future harm as a criterion for current disease could lead to some implausible expansions of the concept, for example, implying that states like pregnancy, which involve increased risks of harm, are diseases (Schwartz 2008, 329-30). And part of the reason that misclassification overdiagnosis occurs is that risk is accepted as evidence of the presence of disease. Further, if we accept that risk is sometimes relevant to disease classification, we again run up against a threshold problem, here the level of risk taken to indicate disease.

Table 1: How influential accounts of disease seem to classify overdiagnoses

Conception of disease	**Misclassification overdiagnosis**	**Maldetection overdiagnosis**
Boorse: BST	Unclear: threshold problem	Yes: dysfunction present therefore there is disease
Nordenfelt: vital goals	Unclear	Unclear: perhaps 'Yes' if early lesions classified as

		disease
Wakefield: HDA	No: Non-harmful therefore no disease	No: Non-harmful therefore no disease
Schwartz: FNC	Perhaps: depending on where boundaries set, and hence risk of future harm	Perhaps: depending upon risk of future harm

Overdiagnosis involves states of physiological dysfunction whose statistical prevalence and potential for harm are not known. The identification of overdiagnosis as a problem strengthens the sense that considering (even statistically abnormal) dysfunction to be sufficient for disease is unsatisfying if it leads to expansions of the concept. Overdiagnosis challenges us to find a more principled way to respond to threshold problems, and to more closely examine the role of risk, or future harm, in the concept 'disease'. To move this discussion forward, we examine the assumptions behind the kind of conceptual analysis typically employed in this literature.

3. The role of philosophical definitions of disease

It is challenging to find criteria for defining disease that provide an extensionally adequate definition; and once identified, such criteria have threshold problems. This observation signifies two prevalent assumptions about our concept 'disease': that it is classically structured, and that it has clear boundaries. In this section, we critique these assumptions and introduce an alternative approach to thinking about philosophical definitions of disease. This alternate approach points to a more fruitful way of thinking about overdiagnosis in relation to disease.

The assumption that disease is a classically structured concept – that it is structured around a set of conditions that are individually necessary and jointly sufficient for its application – lies behind attempts to define disease via a list of such criteria. It is this assumption that validates reasoning from counterexamples as occurs in the traditional debate about defining disease.[17] But this assumption is problematic.[18] Several authors note the continuing failure of philosophers to develop a satisfactory definition that has this structure (Worrall and Worrall 2001, [redacted]). Schwartz (2007b) draws on empirical evidence about concept application and classification which undercuts the idea that concepts, in general, are classically structured. Against the notion that we classify instances with reference to necessary conditions, studies reveal that we typically associate some particulars with specific classes more quickly than others. A robin is classified as a bird more quickly than a goose (Smith et al 1974); and most people will say that more familiar examples of a category, like 'robin' are 'better' examples of that category than others (Margolis and Laurence 1999, 32). Further, our categorisation judgements do not seem to rest on necessary and sufficient conditions. People may say that a creature is a bird because it flies, but also recognise that there are birds that do not fly (Rosch and Mervis cited in Schwartz 2007b, 56).

While there are various possible responses to these problems, they have also led to other ways of thinking about the structure of concepts. On the 'prototype view' of concepts, we classify instances as falling under those concepts, or not, on the basis of their similarity to a prototype. Sadegh-Zadeh (2008), for example, develops an account of disease as a prototype concept, using fuzzy-theoretic logic to define relations of similarity to prototypical disease

[17] See Schwartz (2007b, 50-56), for a detailed demonstration of how this assumption has influenced the kinds of definitions that have been offered, and the way they are developed, in the philosophical debate on disease definition.

[18] A more detailed discussion of these matters is available in [redacted]

states, as a way to avoid these sorts of problems. On the 'cluster' view, a concept is defined in relation to a set of characteristic features – but none of these features is regarded as individually necessary for the application of the concept, while various combinations of the features can be sufficient for application of the concept. The classic example of a cluster concept is 'game'. Wittgenstein (1953) observed the seeming impossibility of finding a criterion necessary to all games, that instead, similarities seemed to 'crop up and disappear' between different games. Thus it is not appropriate to try to define 'game' classically, but we can say that games characteristically involve some of the following: competition, scoring, unchanging rules, be played for entertainment purposes, etc., but any particular game might lack some of these characteristics. [redacted] argue that disease is a cluster concept on the basis that this best fits the heterogenous nature of the category. Characteristic features of disease, such as dysfunction, harm, or increased risk, can be relevant to whether or not we categorise a condition as a disease, even though these criteria may be lacking from some instances.

Thus while there is no agreement about the structure of the concept disease, there is a growing body of work seeking alternatives to a list of necessary and sufficient conditions for defining disease. There seems no reason to retain the prima facie assumption that 'disease' is classically structured.

These alternative accounts also seek to allow that our concept 'disease' does not have clear boundaries; that its borders are vague (Sadegh-Zadeh 2000; [redacted]). Vague concepts are those that have borderline cases which resist classification. Many everyday concepts such 'tall' or 'bald' are vague; that is, it is not clear at precisely what point they apply or fail to apply. Borderline states between tall and not-tall, or

bald and not-bald, resist clarification because what underlies the difficulty with classifying them is that the concepts of baldness and tallness are insufficiently precise to tell us whether or not those cases do or do not fall under it. Insofar as there are many states that are not clearly either disease or non-disease (as well as cases of overdiagnosis, this would include, for example, welcome infertility, obesity, and addiction), it can be understood as a vague concept.

One reason why disease is a vague concept may relate to features of the criteria usually taken to be relevant in defining disease (harm, risk, dysfunction). While the presence of each criterion may influence whether or not some particular condition or case is classed as disease, each may also be present to a greater or lesser extent. 'Harms' range from small inconveniences or minor pains to extreme suffering and death. Increased risk occurs incrementally, such that when we say someone is at increased risk of a consequence, this can mean anything from a low to an extremely high probability of the consequence ensuing. And dysfunction is often a matter of degree. Many dysfunctions lack categorical clarity: a cancer might be more or less invasive or aggressive; the presence of bacteria may be normal at some levels, but indicate disease at higher levels; an infectious disease might be more or less under the control of the immune system.[19] Similarly, the presence of spectral diseases is diagnosed using biomarkers like blood pressure, bone density, or blood glucose, based on diagnostic thresholds that are changeable, contentious, and perhaps not applicable to some cases. In addition to vagueness arising from at least some of the characteristic features of disease being matters of degree, some borderline cases may be so because it is not clear whether they have

[19] One response to this is to consider all of these cases to be disease, and make distinctions about the severity of that disease. Another is to claim that disease itself could be a matter of degree. We do not seek to resolve here which is the preferable response. For a detailed discussion see [redacted]).

enough of these features, or have them in the usual sense.

As a consequence, our concept 'disease' is not determinate enough to clearly apply, or fail to apply to some cases, which will thus be borderline. We can expect borderline cases to occur, in particular, where the characteristic features of the concept are partially present, or where some but not other of the features are present. That is, acknowledging the vagueness of 'disease' leads to the recognition that we may not be able to clarify the disease status of such cases through conceptual analysis based on identifying necessary and sufficient conditions. For the remainder of this paper, we adopt the view that disease is a vague concept that takes the structure of a cluster or prototype, rather than a classical concept with determinate boundaries.

Taking 'disease' to be a vague concept provides a way to avoid choosing between the alternatives of regarding philosophical attempts to define disease as either 'discovery' or 'decision' (Schwartz 2004; 2007b); between deriving a definition from use, and seeking to redefine a concept through argument (Nordby 2007). Instead, we may acknowledge that our existing concept underdetermines its own application, thereby explaining why there are borderline cases which resist classification. But within this borderline zone, philosophical definitions can create further clarity, by revising, in the sense of precisifying, the existing vague concept. The existing concept does not definitively resolve questions of whether controversial or unclear cases – infertility, addiction, obesity, pre-diseases or other risk states, overdiagnosed cases, and so on – are or are not diseases, because it is not precise enough. Where we require greater precision about these borderline cases, we may develop précising definitions. Précising definitions are relative to purpose; that is, to whatever purpose it is that leads us to seek further clarity about the borderline case(s). As such, there may be a range of

different précising definitions of any particular concept, related to different purposes. The précising definition is constrained by, but provides more detail than, the everyday vague concept.

We can now characterise overdiagnosed cases as borderline cases of disease. Further, misclassification overdiagnosis involves threshold vagueness, as cases fall in a borderline zone on spectra of dysfunction, increased risk, and (in some cases) harm. Someone with a borderline level of dysfunction (e.g., blood pressure raised to a level higher than the statistical average), experiencing no current harm, and with some level of risk of future harm; or someone with a slightly raised blood pressure who experiences occasional minor inconveniences as a result, are borderline cases in this sense. Maldetection overdiagnosis involves a different kind of vagueness, since here, there is an identified dysfunction, but whether or not it will lead to harm or indicates increased risk is not clear. Such cases may not have enough of the characteristic features of disease, or lack 'enough' similarity to prototype cases, to be clear cases of disease; or whether they will have enough of the characteristic features may be metaphysically indeterminate, as for example occurs with DCIS or a microscopic papillary thyroid cancer. There may also be cases where both sorts of vagueness, and both sorts of overdiagnosis, are present, such as someone with small to medium sized pulmonary emboli. Whether or not we count such cases as disease will require us to precisify our concept.

4. Defining disease in relation to overdiagnosis

There is a range of possible precisifications that would be consistent with our current concept of disease. In this final section, we stipulatively connect disease status to the imperative to intervene, in order to develop a précising definition of disease that may be of practical use in

addressing the problems of overdiagnosis. To differentiate between the precisified concept and the general concept of disease, we will refer to the former as "disease$_{ODx}$". That is, 'disease$_{ODx}$' will be a more precise concept than 'disease', such that although 'disease' is too vague to tell us whether or not overdiagnosed cases fall under it, we will be able to determine whether overdiagnosed cases fall under disease$_{ODx,}$ and this will guide our practical responses.

What are the features of disease$_{ODx}$? First, as noted above, since the point of having a précising definition is for it to provide practical guidance, disease$_{ODx}$ will involve instances of disease where there is an imperative to intervene. That is, the definition will seek to include cases where it would be a benefit to treat the condition, and exclude those where diagnosis is more likely to harm than benefit. To achieve this, we propose a précising definition with criteria of dysfunction and harm. We need disease$_{ODx,}$ to include cases where there is dysfunction, since this condition prevents the definition applying to cases where someone experiences harm as a result of a range of social factors not usually taken to be indicative of disease (e.g. racial discrimination). In any case, overdiagnosis only occurs when some dysfunction is already identified (on current diagnostic criteria). And we need to include harm, even if we think that disease (simpliciter) could include non-harmful cases, since cases where there is no current or future harm will be cases of overdiagnosis, that we wish to exclude from the category 'disease' for our purposes, through applying our more precise concept of disease$_{ODx}$.

While 'harm' can mean many things, we will usually be able to identify whether a dysfunction is causing current harm (e.g., damage to blood vessels caused by exposure to high blood sugar levels in diabetics). But in cases where there is dysfunction and current

harm, there is no possibility that overdiagnosis will occur: these are cases where overdiagnosis is already ruled out, and can be considered cases of disease (simpliciter). Thus, current harm can be omitted from our definition of $\text{disease}_{\text{ODx}}$, since $\text{disease}_{\text{ODx}}$ only needs to be applied in borderline cases that may or may not be cases of disease.

This leaves us with dysfunction and future harm as necessary for $\text{disease}_{\text{ODx}}$. As discussed in section 1, we will never be able to tell with complete certainty which currently non-harmful dysfunctions will cause harm in the future. It is possible that medical science will develop ways of increasing the accuracy of predictions. For example, in the case of cancer, increasing knowledge of genomics may help us to predict how particular sub-types of cancer will behave in future. But some uncertainty will always remain as a result of metaphysical indeterminacy, as the causal chains leading from a currently identified dysfunction to future harms are open. Thus, in order to talk about future harm we will need to rely on risk estimates drawn from epidemiological studies, as a proxy for knowledge about the future progression in any one individual. Fortunately, there is information available about statistical risk in relation to many health conditions, allowing us to quantify risk based on population estimates. Of course, there are problems with using such data as a proxy for individual risk. Ideally, we should ensure that any statistics used are drawn from a population that is similar in medically relevant ways to the individual to whom they are being applied, and are personalised to the individual's particular conditions to the extent possible. As well as the probability of a harm occurring, the severity of the harm also needs to be taken into account. We propose then the following definition of $\text{disease}_{\text{ODx}}$:

> X is a $\text{disease}_{\text{ODx}}$ iff there is dysfunction that has a significant risk of causing severe harm.

There are a few points worth noting about this definition. First, it takes a classical structure – this is compatible with recognising that disease (simpliciter) does not have such a structure but is vague. Second, the definition is not liable to the problem, noted in section 2, that taking increased risk of harm to be indicative of disease implausibly broadened the concept to include conditions like pregnancy. This potential expansion is prevented by the dysfunction requirement for disease$_{ODx}$. Third, this definition is for use only in borderline cases that are potential cases of overdiagnosis (such as diagnoses of early-stage, non-invasive cancers that are known to be commonly overdiagnosed, or diagnoses of spectral diseases that are close to the diagnostic thresholds and where patients have no other risk factors).[20] Since we have formulated it specifically to deal with such cases, it is not relevant to deciding other kinds of borderline case (like infertility or addiction). Such cases might be decided in relation to some other precisification of the general concept of disease. In order to make decisions about how to precisify a concept, some context or purpose needs to be appealed to, such that these precisifications might look quite different to disease$_{ODx}$. And of course, it would not apply to cases that are already clearly either disease or non-disease, since these do not fall within the borderline zone.

And finally, the definition retains some vagueness, since it leaves open at what level the risk counts as significant, or when a harm counts as severe. We suggest that, within set parameters[21], whether or not a borderline case with dysfunction is disease$_{ODx}$ may be decided individually for different patients. This is because what will count as a significant risk, or a

[20] Though we focus on physical disease, our précising definition may be of value in conceptualising psychiatric overdiagnosis.

[21] These parameters could, for instance include specified kinds or amounts of dysfunction, as well as certain minimum levels of risk. How and by whom these parameters would be set is a further issue that we will not engage with here.

severe harm must be contextualised to the particular patient's other health characteristics, and with reference to their values. Patients differ with regard to what levels of risk they are prepared to tolerate, but also with regard to what outcomes they will count as 'severe'. Some outcomes, for example a massive stroke, are likely to be agreed by everyone to be severe. But others, such as losing a limb or a specific capacity, might take very different values for different people. And even a risk of premature death might be regarded differently by individual patients, for instance by a 20-year-old and an 85-year old. Although this may seem counterintuitive and with potential for significant relativism, contextualising the diagnosis to patients' values occurs only within the predetermined parameters for that condition. In addition, accepting this implication in relation to disease$_{ODx}$ does not imply that it must also be accepted in relation to disease. By stipulating a definition of disease$_{ODx,}$ we are able to provide a definition that helps to clarify the epistemic problems raised by overdiagnosis, and provides a conceptual basis for thinking about potential policy and practice responses in terms of reducing the incidence and harms of overdiagnosis. Such a response is consistent with the claim that disease is not a classically structured concept. The definition of disease$_{ODx}$ tackles the threshold problem by explicitly using qualifiers such as 'significant' and 'serious', acknowledging that these must be further explicated in relation to specific diseases liable to overdiagnosis.

5. Responding to overdiagnosis: is disease$_{ODx}$ useful?

To conclude, we note how our proposal relates to other suggested responses to overdiagnosis. These include strategies such as 'watchful waiting' or using less aggressive treatment options (Barratt 2015); reducing the scope or intensity of screening (Welch et al 2011); encouraging greater informed patient choice about screening (Hersch et al 2013); or requiring new diagnostic tests to be able to identify clinically meaningful levels of disease (Moynihan,

Henry and Moons 2014). Others have suggested revising our systems of disease classification to remove some current diagnostic criteria. This might involve for example, restrictive revisions to the diagnostic criteria for spectral diseases so that less severe cases are classified differently from the more severe. Something like this already occurs with the use of categories 'pre-disease', or 'mild' (though there is debate about the utility of current thresholds for these categories). Moynihan, Henry and Moons (2014) suggest that the harms of overdiagnosis should be a factor in decision-making about diagnostic thresholds; others propose relabelling abnormalities where we do not yet know whether they will be harmful, as 'pseudodisease' or 'IDLE' (indolent lesions of epithelial origin) (Black 1998; Esserman et al 2013; 2014).

All of these responses to overdiagnosis engage with its epistemic features. These authors seek to prevent overdiagnosis or its harms by responding differentially to the information ascribed to a diagnosis; seeking less such information; or re-classifying and re-labelling the information we discover. However, these practical responses fail to engage in the underlying challenge that overdiagnosis poses for our concept of disease.

Our suggestion can be seen as a version of the reclassifying response, recognising some conditions as borderline cases that may or may not be disease, and 'relabelling' them using our definition of $\text{disease}_{\text{ODx}}$. But our argument bases this relabelling on a recognition that the category 'disease' is itself vague. We provide a conceptual grounding for the practically useful claim that potential cases of overdiagnosis should be distinguished from other cases. As well as being intellectually satisfying, this approach provides a response to any potential worry that relabelling strategies are paternalistic.

Recognising that there is a borderline zone rather than a clear distinction between cases that are disease, and cases that are not disease, should encourage both doctors and patients to think differently about cases that fall within this borderline zone. Doing so may be preferable to the alternative of classifying all of these conditions 'disease', but declaring that some of them do not require treatment. In the absence of the distinction provided by $\text{disease}_{\text{ODx}}$, it can be difficult to withhold treatment for diagnosed disease, especially given the current context in which there are strong systemic drivers of overtreatment. In relation to diseases known to be liable to misclassification overdiagnosis, that are usually diagnosed in relation to a threshold point on a spectrum, our discussion implies that we should instead recognise an interval on that spectrum and stipulate that cases falling within that interval may or may not require further attention of various kinds, drawing on both statistical information about the condition, and information about the individual patient. Only after this kind of consideration would the patient be classified with $\text{disease}_{\text{ODx}}$ or no disease. Consider an example of a borderline case: someone with mild hypertension who is experiencing no current harm, but with some level of risk of future harm. On our approach, in order to decide whether this person has $\text{disease}_{\text{ODx}}$, i.e., needs treatment, we must consider what risks their raised blood pressure pose for them, in light of the best statistical information available and their other health characteristics. The absence of other risk factors would count against classifying this as $\text{disease}_{\text{ODx}}$. However, the potential risk of developing more severe hypertension (with attendant risks of stroke, heart attack and so on) are likely to be regarded as severe by most people, which would push the decision back towards counting that case as one of $\text{disease}_{\text{ODx}}$. But the final decision will depend, at least in part, on whether the person would prefer the risks of treatment, or those of non-treatment, and this might come down to personal preferences or other features about the person. For diseases known to be liable to maldetection overdiagnosis, our view would promote recognition that maldetected

abnormalities may or may not really be cases of disease, again referring to risk information to make the best possible decision for individual patients. Someone with an early-stage lesion may well find it easier to accept 'watchful waiting' or less aggressive treatment options if they are told they have a condition that may or may not become a disease, than if they are told they have a disease that may or may not progress. This could help to promote suspension of judgement about treatment, and avoid at least some of the panic and anxiety that can be involved in some treatment decisions.

As noted above, there are various ways to precisify the vague, general concept of disease. One possibility it is worth mentioning here is to count all dysfunctions as disease, but to detach the concept from the imperative to intervene or any other normative implications, as Boorse does (1997). A view along these lines might say that all cases of overdiagnosis are disease, and instead of seeking to define disease in such a way that it includes only cases we want to identify (usually in order to treat), say that we should only treat some cases of disease. Such a route is consistent with our discussion about the concept disease. However, the reason to group together diseases in this way is theoretical: it is a useful way to define the category if our aim is to group together instances that should be studied together. It is a précising definition that is related to a theoretical purpose (as indeed Boorse (1997) acknowledges), and one possible way to precisify the vague concept. But there is no reason to regard it as more authoritative than other precisifications, just because its purpose is theoretical. In seeking to respond to a practical problem like overdiagnosis, other precisifications are more useful.

References

Armstrong SL, Gleitman LR, Gleitman H. 1983. What some concepts might not be. *Cognition* 13: 263-308.

Armstrong D. 1995. The rise of surveillance medicine. *Sociology of Health and Illness* 17 (3): 393-404.

Autier P. 2015. Doubtful health benefit of screening from 40 years of age. *Clinical Oncology*

Barratt A. 2015. Overdiagnosis in mammography screening: a 45-year journey from shadowy idea to acknowledged reality. *BMJ* 350:h867.

Black WC. 1998. Advances in radiology and the real versus apparent effects of early diagnosis. *European Journal of Radiology* 27:116-122.

Boorse C. 1976. Wright on Functions. *Philosophical Review* 85(1):70-86.

Boorse C. 1977. Health as a theoretical concept. *Philosophy of Science* 44:542-573.

Boorse C. 1997. A rebuttal on health. In J Humber & R Almeder (Eds.), *What is disease?*, 3-134. New Jersey: Humana Press.

Carter SM, Rogers W, Heath I, Degeling C, Doust J, Barratt A. 2015. The challenge of overdiagnosis begins with its definition. *BMJ* 350:h869.

Coon ER, Quinonez RA, Moyer VA, Schroeder AR. 2014. Overdiagnosis: How our compulsion for diagnosis may be harming children. *Pediatrics* doi:10.1542/peds.2014-1778.

Engelhardt HT. (1975). The concepts of health and disease. In HT Engelhardt, SF Spicker (Eds.), *Evaluation and Explanation in the Biomedical Sciences*, 125-141. Dordrecht: D Reidel.

Esserman LJ, Thompson IM, Reid B. 2013. Overdiagnosis and overtreatment in cancer: An opportunity for improvement. *Journal of the American Medical Association* 310(8): 797-8.

Esserman LJ, Thompson IM, Reid B, Nelson P, Ransohoff DF, Welch HG, Hwang S, Berry

DA, Kinsler KW, Black WC, Bissell M, Parnes H, Srivastava S. 2014. Addressing overdiagnosis and overtreatment in cancer: a prescription for change. *Lancet Oncology* 15:e234-42.

Glas NA, et al. 2014. Effect of implementation of the mass breast cancer screening programme in older women in the Netherlands: population based study. *BMJ* 349: g5410.

Godlee F. 2012. Preventing overdiagnosis. *BMJ* 344:e3973.

Harach HR, Franssila KO, Wasenius V. 1985. Occult papillary carcinoma of the thyroid: a "normal" finding in Finland. A systematic autopsy study. *Cancer* 56(3):531-538.

Hersch J, Janson J, Barrat A, Irwig L, Houssami N, Howard K, Dhillon H, McCaffery K. 2013. Women's views on overdiagnosis in breast cancer screening: A qualitative study. *BMJ* 346:f158.

Hoda U. 2014. Overdiagnosis of pulmonary emboli by CT pulmonary angiograms. *Thorax* 69:925.

Hofmann B. 2013. Diagnosing overdiagnosis: conceptual challenges and suggested solutions. *European Journal of Epidemiology* 29: 599-604.

Hofmann B. 2016. Medicalization and overdiagnosis: Different but alike. *Medicine, Healthcare and Philosophy* 19(2): 253-64.

Margolis E, Laurence S. 1999. Concepts and Cognitive Science. In E Margolis & S Laurence (Eds.), *Concepts,* 3-81. New Jersey: MIT Press.

Moynihan RN, Cooke GPE, Doust JA, Bero L, Hill S, Glasziou PP. 2013. Expanding disease definitions in guidelines and expert panel ties to industry: A cross-sectional study of common conditions in the United States. *PLOS Medicine* 10(8):e1001500.

Moynihan R, Doust J, Henry D. 2012. Preventing overdiagnosis: how to stop harming the healthy. *BMJ* 344: e3502.

Moynihan R, Henry D, Moons KGM. 2014. Using evidence to combat overdiagnosis and overtreatment: Evaluating treatments, tests, and disease definitions in the Time of Too Much. *PLoS Medicine* 11(7): e1001655.

Nordby, H. 2007. The analytic-synthetic distinction and conceptual analysis of basic health concepts. *Medicine, Health Care and Philosophy* 9:169-180.

Nordenfelt L. 1993. Concepts of health and their consequences for health care. *Theoretical Medicine* 14:277-285.

Nordenfelt L. 2007. Establishing a middle-range position in the theory of health: A reply to my critics. *Medicine, Health Care and Philosophy* 10:29-32.

Pickering TG. 2006. Now we are sick: Labelling and hypertension. *Journal of Clinical Hypertension* 8(1): 57-60.

Rogers WA, Mintzker Y. 2016. Getting clearer on overdiagnosis. *Journal of Evaluation in Clinical Practice* 22(4): 580-7.

[redacted]

Sadegh-Zadeh K. 2000. Fuzzy health, illness, and disease. *Journal of Medicine and Philosophy* 25:605-638.

Sadegh-Zadeh K. 2008. The prototype resemblance theory of disease. *Journal of Medicine and Philosophy* 33:106-139.

Schramme T. 2007. A qualified defence of a naturalist theory of health. *Medicine, Health Care and Philosophy* 10:11-17.

Schwartz PH. 2004. An alternative to conceptual analysis in the function debate. *Monist* 87(1):136-153.

Schwartz PH. 2007a. Defining dysfunction: Natural selection, design, and drawing a line. *Philosophy of Science* 74:364-385.

Schwartz PH. 2007b. Decision and discovery in defining disease. In H Kincaid & J McKitrick (Eds.), *Establishing medical reality,* 47-63. Dordrecht: Springer.

Schwarz PH. 2008. Risk and disease. *Perspectives in Biology and Medicine* 51:320-334.

Smith E, Shoben E, LJ Rips. 1974. Structure and process in semantic memory: A featural model for semantic decisions. *Psychological Review* 81(3):214-241.

Wakefield J. 2014. The biostatistical theory versus the harmful dysfunction analysis, Part 1: Is dysfunction a sufficient condition for medical disorder? *Journal of Medicine & Philosophy* 39:648-682.

[redacted]

Welch HG, Schwartz L, Woloshin S. 2011. *Overdiagnosed: Making people sick in the pursuit of health*. Boston: Beacon Press.

Wittgenstein L. 1953. *Philosophical Investigations,* trans. G.E.M. Anscombe. Oxford: Blackwell.

Worrall J, Worrall J. 2001. Defining disease: Much ado about nothing? *Analecta Husserliana* 72:33-56.