This is a pre-copyedited, author-produced version of an article accepted for publication in the Journal of Medicine and Philosophy following peer review. The version of record: Wendy A. Rogers, Mary Jean Walker, The Line-drawing Problem in Disease Definition, *The Journal of Medicine and Philosophy: A Forum for Bioethics and Philosophy of Medicine*, Volume 42, Issue 4, August 2017, Pages 405–423 is available online at: https://doi.org/10.1093/jmp/jhx010

The published version of this paper is: Rogers W and Walker MJ. The line-drawing problem in disease definition. *Journal of Medicine and Philosophy* 2017; 42 (4): 405-423. DOI: https://doi.org/10.1093/jmp/jhx010.

The line-drawing problem in disease definition

Wendy Rogers and Mary Jean Walker

Abstract
Biological dysfunction is regarded, on many accounts, as necessary and perhaps sufficient for disease. But while disease is conceptualised as all-or-nothing, biological functions often differ by degree. A tension is created by attempting to use a continuous variable as the basis for a categorical definition, raising questions about how we are to pinpoint the boundary between health and disease. This is the line-drawing problem. In this paper we show how the line-drawing problem arises within 'dysfunction-requiring' accounts of disease, such as those of Christopher Boorse and Jerome Wakefield. We then provide several detailed examples to establish that biological dysfunction cannot provide a boundary. We examine potential ways of resolving the line-drawing problem, either by dropping one of the claims that generate it, or by appealing to additional criteria. We argue that two of these options are plausible, and that each of these can be applied with regard to different diseases.

Key words: disease, definition, dysfunction, vagueness, threshold problem

Biological dysfunction is, on many accounts, regarded as a necessary and perhaps sufficient condition for disease. At the same time, disease is conceptualised as either present or absent.[1] This generates a problem. Many biological functions occur across a spectrum, from well functioning to barely functioning. Dysfunction thus admits of degrees, while the presence of disease does not. A tension is created by attempting to use dysfunction, a continuous variable, as the basis for a categorical definition. This raises questions as to whether disease really is all-or-nothing, and if so, how dysfunction can serve as the criterion to pinpoint the boundary between health and disease. This is the line-drawing problem.[2]

In section one, we examine how the line-drawing problem arises for accounts of disease that

[1] We refer to this conceptualisation of disease as being either present or absent as "all-or-nothing". This refers to the boundary between a disease state and a non-disease state, which we take to be conceived of as categorical. On categorical accounts, once present, disease states may vary in severity, so the "all" refers to the presence of disease, rather than its degree.
[2] The phrase is from Peter Schwartz's (2007a) work, to which we are indebted, though Schwartz frames the problem in a different way to us.

use dysfunction as a criterion, such as those of Christopher Boorse and Jerome Wakefield. In section two, we utilise several detailed examples to establish that biological dysfunction cannot provide a determinate boundary. We then, in section three, examine potential ways of resolving the line-drawing problem. We argue that two of these options, namely dropping the assumption that disease is all-or-nothing, or recognising that some borderline cases will be indeterminate, are plausible. We suggest that the line-drawing problem might be resolved differently in different cases. Specifically, we argue that for diseases that are defined on the basis of variables that are measurable on a spectrum, we should drop the assumption that disease is all-or-nothing. But for diseases where the spectrum of conditions between an unambiguous disease state and a clearly healthy state is relatively circumscribed, such as some cancers, it can be useful to retain the binary conception of health and disease, while recognising that the boundary between them is vague, so that there may be 'borderline' cases that are not obviously either disease or non-disease.

1. Health, disease and the dysfunction requirement

Many influential approaches to defining disease assume that disease and health are mutually exclusive (Sadegh-Zadeh 2000, 606). If one has a disease, one is not healthy, and healthy individuals do not have any diseases. Boorse explicitly describes his project as analysing "the notion of *disease* behind the view that health is the absence of disease" (1977, 550, italics in original). Elsewhere he states "'disease' is simply synonymous with 'unhealthy condition'" (Boorse 1975, 50) and equates theoretical health with being free of disease (1975, 57).[3] Most contributors to the debate about defining disease adopt a similar assumption, and follow Boorse in regarding health and disease to be mutually exclusive. Wakefield, who uses the

[3] In 1987, Boorse adopted the phrase 'pathological condition' in place of 'disease' as better reflecting the broad notion he aims to capture, stating that the aim of his theory is to analyse the distinction between the normal and the pathological (1987, 7). We use 'disease' below in this broader sense.

The published version of this paper is: Rogers W and Walker MJ. The line-drawing problem in disease definition. *Journal of Medicine and Philosophy* 2017; 42 (4): 405-423. DOI: https://doi.org/10.1093/jmp/jhx010.

term 'disorder' rather than 'disease', indicates that this concept serves "a similar purpose" to Boorse's use of 'disease' or 'pathology' (2014, 653). At the same time, many accounts include biological dysfunction in their analysis of disease. A tension is generated by the differing logical forms of the function/dysfunction spectrum and the health/disease dichotomy. In what follows, we develop this point by looking closely at how dysfunction is conceptualised in the accounts of Boorse and Wakefield.

Boorse's biostatistical theory (BST) defines disease as "a type of internal state which impairs health, i.e., reduces one or more functional abilities below typical efficiency" (Boorse 1977, 562). 'Function' is defined biologically in terms of causal contributions to a goal. That is, the function of a cell, organ, or system is its species-typical contribution to a biological goal. 'Normal' functioning is "performance by each internal part of all its statistically typical functions with at least statistically typical efficiency, i.e. at efficiency levels within or above some chosen central region of their population distribution" (Boorse 1977, 558-9). Thus a normally functioning heart in a 30-year-old woman pumps blood at or above the average efficiency for her reference class (i.e. taking account of sex and age). Dysfunction, then, occurs when functional efficiency falls some distance below the population mean (1977, 589).

Boorse addresses the line-drawing problem by appealing to the statistical frequency of a dysfunction, stipulating that abnormal function is that falling more than a certain distance below the population mean. He goes on to say:

> [T]his distance can only be conventionally chosen, as in any application of statistical normality to a continuous distribution. The precise line between health and disease is usually academic, since most diseases involve functional deficits that are unusual by

The published version of this paper is: Rogers W and Walker MJ. The line-drawing problem in disease definition. *Journal of Medicine and Philosophy* 2017; 42 (4): 405-423. DOI: https://doi.org/10.1093/jmp/jhx010.

any reasonable standard (1977, 559).

In discussing reference ranges in medicine, which use conventional cut-offs for defining 'normal' such that abnormal results are regarded as those falling more than two standard deviations from the mean, Boorse notes that the use of two standard deviations is arbitrary given the vague boundaries of pathological states (1987, 371, cited in Schwartz 2007, 374). Boorse thus recognises that biological functions have continuous distributions and that pathological states can have vague boundaries. But despite these observations, he claims that in most, or almost all, cases there will be an obvious level of dysfunction that marks out disease from health (Boorse 1997; 2014). It is, however, unclear as to what feature some level of dysfunction should have that makes it pathological rather than normal, especially if a statistical definition of abnormality is recognised to be arbitrary. Presumably Boorse would not want to appeal to some feature not included in the definition of the BST in his reference to "any reasonable standard". Rather, his point appears to be that this question is unimportant or "merely academic", perhaps because he regards it as a problem only in a minority of cases. In contrast, we think that many more cases pose this problem than are typically assumed, as we seek to show in section 2. Even paradigmatically 'all-or-nothing' diseases like cancer or infections may sometimes have a vague boundary with health.[4]

Schwartz (2007a, 375-6) argues that statistical frequency cannot be used to solve the line-drawing problem, as disease might be quite common in a population, or it might be very rare or absent in a population. Where disease is common, this approach to drawing the line will imply that those we would intuitively recognise as diseased actually are not; where it is rare,

[4] One possibility for the BST would be to say that diseases are such in virtue of the total dysfunction of a very small part. This could provide a way to draw the line by including all cases where a small part is dysfunctional. We do not dispute that drawing the line in such a way is theoretically consistent, nor that if the project is to draw some line then drawing it wherever there is any pathology, no matter how small or unimportant, might be as good a place as any, and perhaps even the most useful place theoretically speaking. But choosing such a location needs to be argued for and not assumed; and it disadvantages us if it prevents recognition of degrees of disease.

the approach would imply that people we would normally think of as healthy are diseased. The frequency approach has implications that do not match our usual uses of these terms. This critique makes clear the oddity of claiming that a state's statistical frequency – a feature that is extrinsic to any particular level of function or dysfunction – is determinative as to whether any particular instance of that state counts as disease. If statistical frequency is determinative in the way Boorse seems to suggest, then whether or not I have a disease depends on facts about other people, rather than just on facts about my own condition. Although it would be possible to accept extrinsic criteria such as frequency despite this oddness, Schwartz notes that it is the effects a state has on an organism, not the state's statistical frequency, that actually seem to be relevant in deciding disease status (Schwartz 2007a, 376-7).

Wakefield's use of the concept of dysfunction in defining disease differs from Boorse's in that he does not take it to be sufficient for disease, and he uses an aetiological account of function. On Wakefield's harmful dysfunction analysis (HDA), conditions are diseases[5] if and only if: a) the condition fails to meet the value criterion (the details of which do not concern us here); and b) "the condition results from the inability of some internal mechanism to perform its natural function" (1992a, 384). Natural functions are "effects that explain the existence and structure of naturally occurring physical and mental mechanisms" (1992a, 383); that is, the function of the eyes is to see, as the fact that eyes facilitate vision helps to explain why we have eyes (1992a, 383). Dysfunction occurs when a mechanism fails to perform its natural function (1992a, 383) and "exists when a person's internal mechanisms are not able to function in the range of environments for which they were designed" (1992b,

[5] Though Wakefield uses the term 'disorder' rather than 'disease', we do not think anything significant follows from our use of 'disease' throughout to refer to conditions identified as disorders on the HDA.

The published version of this paper is: Rogers W and Walker MJ. The line-drawing problem in disease definition. *Journal of Medicine and Philosophy* 2017; 42 (4): 405-423. DOI: https://doi.org/10.1093/jmp/jhx010.

243). Again, the function of the 30-year-old woman's heart is to pump blood, and she has a healthy heart insofar as it pumps blood in a way selected for by evolution (as nature 'intended').

The appeal to natural selection seems to provide little guidance regarding the level of functionality at which we should say that a bodily system, organ, tissue, or cell is no longer performing in the way for which it was selected. Schwartz (2007a) argues that appealing to natural selection cannot help us determine a level of function to be considered healthy because what is selected for or against depends, not just upon how well a particular organ functions, but also on other variations existing in the population, and how these compare (2007a, 370-1). However, as the HDA also has a value criterion, Wakefield does not require the notion of dysfunction to bear the whole weight of defining disease. He explicitly links the level of dysfunction to harm in order to distinguish between diseases and conditions that are abnormal but not disease: "a part dysfunction/lesion [is] a disorder only if the deviation in the functioning of the part affects the well-being of the overall organism in a harmful way" (1992a, 375). He does note, "discovering what in fact is natural or dysfunctional may be extraordinarily difficult" (1992b, 236), but does not investigate the implications of this difficulty for his theory.

Thus, both the BST and the HDA seem to require a way to determine decisively when dysfunction occurs, in order to demarcate between health and disease. This will be determined statistically for Boorse, or via appeal to natural selection for Wakefield. But there are problems with drawing this line by appeal either to statistics or natural selection. Following his detailed engagement with these issues, Schwartz (2007a) proposes an alternative model which he calls the 'function and negative consequences' (FNC) view. The

The published version of this paper is: Rogers W and Walker MJ. The line-drawing problem in disease definition. *Journal of Medicine and Philosophy* 2017; 42 (4): 405-423. DOI: https://doi.org/10.1093/jmp/jhx010.

FNC view adds a 'negative consequences' requirement to Boorse's statistical approach, so that a dysfunction is a disease only when it has negative consequences as well as falling some distance below a mean. The negative consequences are cashed out in terms of effects that "significantly diminish the ability of a part or a process in the organism, or of the overall organism, to carry out an activity that is generally standard in the species and has been for a long period of time" (2007a, 379). It seems unclear how this view will succeed at drawing the line where the BST and HDA have faltered, since neither function nor 'negative consequences' provide a clear line to which to appeal. Schwartz does note that his method "provides no simple rule about where to put the line", instead clarifying how the addition of negative consequences allows responses to the problems facing the frequency approach (2007a, 377). Schwartz's approach relies heavily upon identifying "significant decreases" (2007a, 380) in function, and he writes: "Admittedly, there is some ambiguity how to define "significant" in these settings, but this may be unavoidable" (2007a, 383). Schwartz thus acknowledges that his solution is also subject to problems concerning line-drawing, suggesting this may be insoluble. More importantly for our purposes, he does not attribute the line-drawing problem to the attempt to base a categorical definition on a continuous variable, or consider it to threaten the binary characterisation of health and disease.

All of these authors, Boorse, Wakefield and Schwartz, hold that disease is either present or absent, denoted by the presence or absence of dysfunction (with or without additional criteria). While there is some admission that it is challenging to identify the level of dysfunction that marks the boundary of disease, this issue is largely disregarded on the assumption that, by and large, dysfunction can be readily identified in contrast to adequate or healthy functioning in the majority of cases. Yet views about the categorical nature of disease rely upon the robustness of this assumption. In the next section we examine dysfunction

The published version of this paper is: Rogers W and Walker MJ. The line-drawing problem in disease definition. *Journal of Medicine and Philosophy* 2017; 42 (4): 405-423. DOI: https://doi.org/10.1093/jmp/jhx010.

across three disease types to establish that, in a range of conditions, function and dysfunction occur on a spectrum.

2. Can dysfunction provide a categorical boundary?

Insofar as the line-drawing problem has been discussed, this has largely been with regard to disorders that are defined in terms of explicit thresholds on continuous biological variables (Schwartz 2007a; 2008). Hypertension (Stage 1), for example, is defined as blood pressure above 140/90 mmHg (Joint National Committee on Prevention, Detection, Evaluation, and Treatment of High Blood Pressure 2003); type 2 diabetes as $HbA_{1c} \geq 6.5$ (or fasting plasma glucose ≥7 mmol/L or 2 hour plasma glucose ≥ 11 mmol/L following oral glucose tolerance test) (American Diabetes Association 2015); obesity as a body mass index greater than or equal to 30 (American Medical Association House of Delegates 2013); and so forth. The boundaries defining these conditions have shifted over time, generally to create more inclusive disease categories (Moynihan et al 2013).

The changing nature of these thresholds suggests that there is no incontrovertible point at which function switches to dysfunction. There is often continuity, with severe disease at higher levels of the relevant variable (indicated by increasing mortality and morbidity at, for example, extremely elevated blood sugar or blood pressure levels), while lower levels of the relevant variables are classified as normal or healthy. For each condition, the threshold could plausibly be higher or lower by at least one increment of the relevant measure without greatly altering the nature of the disease (e.g., hypertension could plausibly be 139/89 or 141/91 mmHg). The increments that separate having a disease and not having it are so small that the exact position of the line demarcating disease from non-disease – the diagnostic threshold – seems arbitrary. These conditions appear to be matters of degree, tracking the trajectory from

The published version of this paper is: Rogers W and Walker MJ. The line-drawing problem in disease definition. *Journal of Medicine and Philosophy* 2017; 42 (4): 405-423. DOI: https://doi.org/10.1093/jmp/jhx010.

function to dysfunction; from health to disease. The lack of a clear boundary is also indicated by the way that these conditions morph seamlessly into risk factors (pre-hypertension, pre-diabetes, overweight, etc.), suggesting that there are only shades of grey rather than quantifiable functional deficits between, for example, healthy blood pressure, pre-hypertension, and overt hypertension (see, e.g., Pickering 1959). Thus it seems implausible to claim that the presence of dysfunction can distinguish between health and disease regarding these disorders defined in terms of continuous variables and their associated risks.

The preceding conditions might be considered fairly obvious examples in terms of criticising dysfunction-requiring accounts, because they are diseases that are defined in relation to variables occurring on a standard distribution curve. This feature immediately raises the question as to where to draw the line between healthy and pathological given the lack of disjunctions. The lack of a clear demarcation is reflected in the actions of the expert committees that determine the boundaries of these diseases, as they explicitly take account of practical considerations rather than claiming to identify the point where pathology starts. This may be a purely pragmatic approach, but seems to support the view that there is no clear boundary to find.[6] Let us consider, then, the boundary between function and dysfunction in some apparently more categorical diseases, such as cancer or infections.

The identification of cancer used to be relatively straightforward. Prior to the advent of screening programs, most cancers were diagnosed because patients noticed something wrong. People presented with symptoms related to the abnormal growth of cells or tissues, such as a hard lump in a breast, back pain from multiple myeloma, or headaches and blurred vision from a brain tumour. Lesions were often clinically visible or palpable, or detectable with

[6] See, e.g., Hofmann (2016) for a discussion of the pragmatic arguments for designating obesity a disease.

relatively unsophisticated tests. The categorical nature of disease was seemingly evident in these cases: patients either had, or did not have, cancer. There were fewer grey areas, because any abnormal cell growth had declared itself malignant by the time the patient noticed something wrong and offered him or herself for diagnosis.

However, changes in our understanding of cancer, and in our technologies for identifying changes in cell growth and predispositions to such changes, have seriously undermined the diagnosis of cancer. The abnormality of function evident with invasive cancerous growth and distant spread is no longer so distinct with the early identification of atypical or dysplastic growth. Now it can be challenging to declare any particular cell or cells abnormal or malignant, as there are many stages in between, where function starts to go awry in almost imperceptible ways.

The classification of neoplasms or tumours is complex. Initial classifications were topological. Tumours were largely defined in relation to their anatomical location: breast, prostate, lung and so forth (ICD-6 cited in Muir and Percy 1991). This was rapidly supplemented with morphologic (or histologic) criteria relating to the appearance and architecture of cells within the tumour, and criteria relating to behaviour. Tumours with characteristic architecture and malignant (i.e. invasive) behaviour are called cancers (Muir and Percy 1991). Cancer cells grow in a disorderly, chaotic manner, unresponsive to local requirements or regulation, leading to invasion of local tissues and spread to other parts of the body (known as metastases). This tripartite classification (topology, morphology, behaviour) was disrupted by the advent of genetics, with the first cancer-specific genetic abnormality reported in 1960 (Brandal et al 2010). Currently it is thought that genetic aberrations are the central pathogenic event in neoplastic transformation (the somatic mutation theory of cancer).

The published version of this paper is: Rogers W and Walker MJ. The line-drawing problem in disease definition. *Journal of Medicine and Philosophy* 2017; 42 (4): 405-423. DOI: https://doi.org/10.1093/jmp/jhx010.

Various genetic alterations in cell physiology have been identified that collectively produce cancerous cells characterised by a lack of the usual growth-regulatory mechanisms (Soto and Sonnenschein 2004). Cancer develops and progresses by gradually acquiring new mutations and other changes that allow cancer cells to expand and migrate (Greaves and Maley 2012).

The main point is that as our understanding of cancer becomes more sophisticated, it becomes increasingly difficult to pinpoint the relevant dysfunction, and hence determine where any line should be drawn between the presence or absence of cancer. If at cellular/tissue level, where the pathognomic dysfunction is local invasion and/or distant spread (Lazembnik 2010), then we cannot confidently diagnose cancer until one of these events has occurred. This may or may not provide us with a categorical definition, depending upon how possible it is to establish local invasion and metastatic spread. On the other hand, if we rely upon genetic markers to identify dysfunction at the level of the deregulated cell growth pathways, we may not track cancer, as not all cells with the relevant genetic markers will proceed to malignant growth.

Whether we take a morphologic or genetic approach to defining cancer, it seems that cancerous growth occurs on a spectrum of deregulation. There may be cells with genetic markers but no subsequent abnormalities in growth, abnormal but non-malignant cell growth, malignant-looking tumours that do not metastasise, all the way through to cancer with distant metastases. For at least some neoplasms (e.g. cervical, oesophageal, breast, prostate)[7] there is no distinctive point at which we can say that the dysfunction underlying the disease of cancer is, in Boorse's terms, "unusual by any reasonable standard": dysplasia occurs on a spectrum.

[7] There are clearer boundaries in some neoplasms such as the leukaemias, in which abnormal cells are either present or absent. But for many of the endothelial neoplasms, there is a spectrum of disordered growth.

The published version of this paper is: Rogers W and Walker MJ. The line-drawing problem in disease definition. *Journal of Medicine and Philosophy* 2017; 42 (4): 405-423. DOI: https://doi.org/10.1093/jmp/jhx010.

Like the chronic diseases discussed earlier, it is difficult to find a demarcation between normal function and dysfunction. At one end of the scale, cancers are unmistakably present, but at the other, there are atypical cells that may or may not proceed on a malignant trajectory. And we cannot pinpoint any one mutation among the many acquired by the clonal cells that necessarily determines a malignant trajectory, and hence the presence of cancer.

A Boorsian might respond by claiming that the disease cancer begins at a particular point, for example, with the very first abnormal cell growth, i.e. mild dysplasia. But this seems arbitrary in that whatever dysfunction occurs at a histological level is a reflection of sub-cellular dysfunction, where genetically-programmed cellular processes go awry. Searching for the dysfunction that marks health from disease in the case of cancer triggers an infinite regress down to molecular variations, none of which might be the sole cancer-causing dysfunction. It seems there is no specific dysfunction that picks out cancer. Not all genetic mutations lead to abnormal cell growth; not all cases of atypia progress to dysplasia, and so forth. Any dysfunction that we take to be pathognomic of cancer is going to be arbitrary as some other part- or process-dysfunction could equally well be selected.

Our third example involves diseases caused by infectious agents. The presence of diseases like the measles, small pox, syphilis, meningitis, septicaemia and so forth seems incontrovertible. There are typical signs and symptoms; there may be positive identification of relevant organism; rises in relevant immune titres provide serologic evidence; responses to treatment may be characteristic; and so forth. All of these make it possible to say that a person definitely has the measles, meningitis, or syphilis.

Not all infections, however, admit of such certainty by establishing themselves in a clear cut

way. Consider infection of the bladder with the common pathogen *Escherichia coli*, which is associated with diseases including asymptomatic bacteriuria and acute cystitis (also known as acute uncomplicated urinary tract infection – UTI). UTI is characterised by symptoms of urgency, frequency, dysuria (pain on voiding), and haematuria (blood in urine) (Schnarr and Smaill 2008; Giesen et al 2010). The gold standard for diagnosing UTI is detection of a pathogenic organism in the urine, in a symptomatic person. The diagnostic threshold is quantitative and measured by the number of colony-forming units (cfu) of bacteria cultured from a specimen of urine. However, there is no consensus as to what level of bacteriuria is required in order to establish the presence of disease (Schmiemann et al 2010).

The example of UTI points to two challenges in drawing the line between a healthy lower urinary tract and one with infection. The first relates to function. For disease, we need to identify dysfunction, i.e., what function it is that is failing. The presence of bacteria in the urine would seem to indicate that some protective function is not working properly given that the urinary tract is usually sterile (Macejko and Schaeffer 2007). However, the transient presence of at least some *E. coli* in the urine is normal in many individuals and is not associated with any symptoms of UTI such as discomfort, urgency, pain on voiding, suprapubic pain and so forth. Perhaps the relevant dysfunction is when *E. coli* overcomes the body's protective mechanisms to the extent that bacteria successfully multiply within the bladder. But this could be the result of various factors including incomplete bladder emptying (as occurs for example in outflow obstruction secondary to enlargement of the prostate); high levels of sugar in the urine (in a person with poorly controlled diabetes); pregnancy-related changes such as relaxed smooth muscle leading to decreased peristalsis of the ureters and stasis in the bladder (Macejko and Schaeffer 2007); or some failure of the immune system. This presents a bit of a puzzle, as before we can draw a line between healthy and unhealthy

function, we need to know what the relevant function is, which seems difficult to identify in the case of UTI. Proponents of dysfunction-requiring accounts might respond that all of these are dysfunctions (with the possible exception of pregnancy-related changes), and that the presence of any one of them is pathological. But this does not necessarily solve the problem, as in each case it would be possible to have the pathological change (sugar in urine, incomplete emptying etc.) without having a UTI.

Setting this issue aside, the second challenge relates to the diagnostic threshold for UTI: how many cfu/mL demonstrate the relevant dysfunction? Quantitative thresholds are used to define UTI. Accepted diagnostic thresholds range between bacterial counts of $\geq 10^2$ cfu/mL to counts of $\geq 10^5$ cfu/mL (Giesen et al 2010). This nomenclature implies that there may be more or less than 10^5 cfu/mL of *E. coli* cultured from an individual's urine – the presence of bacteria occurs on a spectrum from zero to $\geq 10^5$ cfu/mL. As with any continuous variable, the choice of a specific threshold is to some extent stipulative or conventional, rather than tracking a clear change in function. The bladder's protective mechanisms may be more or less effective at keeping the urine clear of bacteria, and at what point, if any, we declare a certain level of cfu/mL to be pathogenic is a matter of choice rather than discovery. Thus with UTI, the relevant dysfunction is difficult to specify, and one of the strongest indicators of dysfunction, the bacterial count, is conventionally chosen, contested, and may be present with or without UTI (Schmiemann et al 2010).

We find a similarly complex picture with other infections such as TB. A recent review (Lawn et al 2011) suggests that the conventional understanding of latent TB and active TB as two distinct conditions is erroneous. Rather, they argue, the host-pathogen relationship that occurs between humans and *Mycobacterium tuberculosis* occurs on a spectrum. At one end,

infection is eliminated by the innate immune system, while at the other, there is uncontrolled bacterial replication and clinical TB disease. In between there are stages: the infection may be eliminated or remain under immune control; or active bacterial replication may be subclinical (see Lawn et al 2011, Table 2, p3).

Immune control features heavily in descriptions of TB, making it plausible to consider that TB is a disease characterised by a failure of the immune system to eradicate or control infection by *M. tuberculosis*. But given the various interactions and stages of control that are possible between the host's immune system and the pathogen, it is not clear which stage signifies the presence of dysfunction, and hence disease. Convention decrees that latent TB (corresponding to the infection being under control, or subclinical bacterial replication) is not a disease while active clinical TB is (Lawn et al 2010). It might be argued that all of the stages of TB infection are pathological, even if some are not associated with clinical disease. But if Lawn et al (2010) are correct in their claim that we should think of the host-pathogen relationship as a spectrum, we are again faced with the question of where on the spectrum dysfunction occurs. We can imagine cases where individuals are exposed to TB but their immune system eradicates the bacilli that manage to enter the respiratory system, and this successful response might take more or less time. Do they have an infection during the time it takes to eliminate the bacilli, and does this mean we should identify dysfunction as the triggering of any part of the immune response, no matter how swift and successful?

With these detailed examples, we hope to have shown that, for a range of disparate conditions, dysfunction cannot provide a categorical boundary between disease and non-disease. For chronic 'lifestyle'-related diseases, for some malignancies, and for at least some infections, the relevant dysfunctions admit of variety: cell growth may be more or less

The published version of this paper is: Rogers W and Walker MJ. The line-drawing problem in disease definition. *Journal of Medicine and Philosophy* 2017; 42 (4): 405-423. DOI: https://doi.org/10.1093/jmp/jhx010.

disordered; the immune system may be more or less in control of invading pathogens; and physiological parameters may be nearer to or further from healthy. At one end of each spectrum there is overt disease with obvious dysfunction, while elsewhere, there is no disease and normal function is achieved, with every shade of grey in between. And this makes intuitive sense: biological functions seem to be the kind of things that may perform better or worse, in increments across the full range. Biological functions may categorically cease altogether (the heart may stop beating, the liver stop metabolising and the kidneys stop filtering blood), but short of absolute cessation of function, there are degrees of performance all the way up to abundantly healthy levels.

3. Options for resolving the line-drawing problem

If we accept that dysfunction occurs by degree, the claim that disease is analysable in terms of dysfunction is in tension with the assumption that disease is either present or absent, all-or-nothing. Accordingly, we have four options:

1. Retain the assumption that disease is all-or-nothing, and abandon attempts to analyse it in terms of dysfunction;
2. Drop the assumption that disease is all-or-nothing, and retain the view that it is analysable in terms of dysfunction;
3. Retain both claims and investigate other methods for resolving the tension between them; or
4. Drop both claims. As this option seems irrelevant to our interests, we will say no more about it.

	Disease is analysable in terms of dysfunction	
	Yes	No

The published version of this paper is: Rogers W and Walker MJ. The line-drawing problem in disease definition. *Journal of Medicine and Philosophy* 2017; 42 (4): 405-423. DOI: https://doi.org/10.1093/jmp/jhx010.

Disease is all or nothing	Yes	3)	1)
	No	2)	4)

Option 1

The tension we have identified could be regarded as a reason to prefer an analysis of disease that does not appeal to dysfunction. Alternative accounts of disease might, for example, appeal instead to harm, disruption of normal activities, or increased risk of harm or premature mortality. This is unsatisfactory if we agree with the widely shared notion that dysfunction does capture, or partly capture, what we mean when we call something a disease – and many do seem to have this intuition (cf. Hausman 2012, 524). The notion of dysfunction seems to capture efficiently the sense that disease involves something amiss with our physical or psychological states, and this intuition is potentially consistent with both naturalist and normativist approaches. Next, dropping dysfunction from an analysis of disease challenges us to find an alternative account that can actually resolve the line-drawing problem, and rival accounts do not clearly do so. Harm, increased risk of harm, or disruption of normal activities, all seem to involve degrees. This indicates that dropping dysfunction from our analysis will not necessarily help to resolve the problem; perhaps we will face a similar problem on many, or even any, analyses of disease. Thus although dropping dysfunction from our analysis of disease would be consistent with our argument so far, we do not consider it a very satisfactory solution as it seems to shift rather than solve the problem. Accordingly, the rest of our discussion will focus on options that retain the claim about dysfunction. We note however, that anyone wanting to drop dysfunction from the analysis of disease on other grounds might regard our argument as an additional reason to do so.

The published version of this paper is: Rogers W and Walker MJ. The line-drawing problem in disease definition. *Journal of Medicine and Philosophy* 2017; 42 (4): 405-423. DOI: https://doi.org/10.1093/jmp/jhx010.

Option 2

If we do wish to retain dysfunction in the analysis of disease, we might instead abandon the claim that disease is all-or-nothing (i.e. having a determinate threshold), and instead regard disease as a matter of degree. The degree of disease present in any particular case might be thought to directly track degrees of dysfunction or, if we think the concept has other features, to track degrees of dysfunction as well as other criteria that are present to varying extents. This kind of view would be consistent with some of the ways we talk about disease and health, for example, when we say someone is more or less healthy. It may also be consistent with some of our current categorisations that recognise degrees of disease: we recognise risk factors or pre-diseases that can worsen into diseases; and stages of disease such as cancer that reflect levels of severity. A 'continuum' approach to defining disease can also accommodate cases where dysfunction does seem to be categorically either present or absent, of course, as such cases fall at the extremes of the spectrum.[8]

Option 2 responds to the line-drawing problem, then, by saying that there isn't any line to draw. A view consistent with this option is developed by Sadegh-Zadeh (2000), who uses a fuzzy-theoretic approach to develop a way to represent degrees of disease (discussed further below). Hausman's functional efficiency theory would also, we take it, be compatible with this option (2012, 2014). Hausman argues that we cannot distinguish between health and disease on the basis of statistical frequency, and that while we do need ways to understand the levels of functional efficiency of bodily systems or parts, which levels are called 'healthy'

[8] It might be argued that we are conflating 'degrees of disease' with 'degrees of severity' of disease; that where any level of disease is present, we should simply say that disease is present, and that its degree of severity is then a further specification to be made. We reject this distinction because it seems to assume, rather than establish, that there is a clear distinction to be drawn between disease and health. But where diseased and healthy states fade gradually into each other – where the boundary between them is vague – how we are to determine such a line, and whether we should even think in terms of drawing lines, is exactly what is under dispute.

and which 'pathological' is not important (2012, 534).

As Hausman recognises, this does not preclude the drawing of lines for various practical reasons. For instance, we may need to decide whether or not to treat a condition, whether or not to excuse someone from work, whether someone is entitled to disability benefits, and so on. Such 'all-or-nothing' decisions are often based on information that is a matter of degree, such as estimates of risk[9], so this is not necessarily problematic. Practical decisions that require an 'all-or-nothing' kind of response can also be dealt with by appealing to stipulative lines. Indeed, something like this already occurs. Even if the philosophical discussion has assumed that disease is all-or-nothing, medical practice and everyday use of the terms recognise that there are borderline cases between rude health and obvious disease, and draw boundaries as necessary for decision-making or other practical purposes. These boundary-drawing practices might use a range of criteria, including practical and evaluative criteria as well as measures of dysfunction. Something like this seems to occur, for example, when expert committees or consensus panels decide on where to set diagnostic boundaries for specific diseases. These panels may consider factors like attempting to remedy underdiagnosis, enabling earlier diagnosis, or facilitating research or management, in deciding how to draw these stipulative lines (Moynihan et al 2013). Similarly, practitioners may use risk estimates, and / or clinical judgement, to decide on how to draw a stipulative line in individual cases.[10] Option 2 is consistent with considering 'disease' itself to be a

[9] Risk estimates occur by degrees because they reflect probabilities, not because they reflect that there are degrees of disease, but this demonstrates that treatment decisions may rely on criteria other than dysfunction in any case. See Vickers et al (2008) for an argument that risk estimates are a better basis for treatment decisions than diagnosis, since diagnosis is also binary.

[10] Another option here may be to alter the consequences attached to the practical decisions such that the decisions do not need to be 'all-or-nothing' and can themselves respond to the presence of degrees. For instance, a system in which physicians issue a note stating what work a person is capable of doing, instead of simply signing a person off work entirely for a period has recently been instituted in the UK (UK Government 2015). Of course, this might not be possible for all practical decision-making needs.

matter of degree, while appealing to a conventional or stipulative line where appropriate.[11] Option 2, retaining dysfunction as a criterion in defining disease, but accepting that whether or not disease is present can be a matter of degree, seems therefore quite plausible.

Option 3

There may be some ways to retain both the claim that disease is all-or-nothing, and the claim that it can be analysed in terms of dysfunction, if there are other ways to resolve the tension which arises from the different logical forms of disease and dysfunction. One way to do this would be to include additional criteria along with dysfunction in an analysis of disease, and argue that while dysfunction cannot definitively guide us in drawing the line, other criteria that are part of our concept of disease can so guide us.

This option is potentially consistent with Wakefield's HDA, or Schwartz's FNC view. Both views accept that dysfunction is part of our concept of disease, but regard other criteria – harm or negative consequences – as also relevant. As we saw above, and as both authors recognise, neither of these views necessarily tell us exactly where to draw the line (see Wakefield 1992b, 236; Schwartz 2007, 377). Indeed, as noted previously, many of the other criteria often brought up in relation to disease, such as harm, or increased risk of adverse consequences, are also matters of degree. Thus, while it is possible that some other criteria, or combination of criteria, might enable us to identify disease boundaries somewhere on the spectrum of dysfunction, it appears that finding them is, at least, highly challenging.

[11] There are also other situations where it may be helpful to consider disease to be simply present or absent, such as for pedagogical purposes in medical education, or when undertaking differential diagnosis. Again, we can here consider the 'line' to be stipulative and playing a practical role.

The published version of this paper is: Rogers W and Walker MJ. The line-drawing problem in disease definition. *Journal of Medicine and Philosophy* 2017; 42 (4): 405-423. DOI: https://doi.org/10.1093/jmp/jhx010.

Another option, should we want to retain both claims, is afforded by recognising vagueness in our concept of disease. Vague concepts are those that have borderline cases that resist clarification, exemplified by sorites problems. Sorites problems involve gradual change from one state to another which reveals the lack of a clear location of the boundaries of our concepts. For instance, the point at which a growing number of grains of sand count as a 'heap', how much hair loss is needed for a person to count as 'bald', or the exact height at which someone counts as 'tall' are all vague. Our inability to draw a line between health and disease could be an instance of problems often encountered in relation to vague terms. There are several possible responses to this problem of vagueness, such as accepting that concepts can apply by degree; believing that there is a definite border but that we cannot identify it; or accepting that, for some concepts, there are cases that are stubbornly indeterminate.

Recognising that a concept is vague means recognising that it has borderline cases which only partially meet that concept's criteria. Vagueness in a concept can be understood to reflect that, at the borders, concepts can be a matter of degree. For example, regarding the concept 'disease' and the claim that dysfunction is required for disease, there seem to be borderline cases of dysfunction where it is true that a degree of dysfunction is present, but it is also true that there is a degree of normal function. In these cases, it is neither wholly true nor wholly false that the dysfunction is a case of disease. This response to vagueness is formalised in what is known as many-valued logic (Sorensen 2012), which allows that propositions can have any degree of truth between 0 (false) and 1 (true). On this kind of view, a proposition like "X is a disease" can take any truth value between 0 and 1 in particular cases. Thus "X is a disease" might be, for instance, 0.5 true or 0.2 true in particular cases, which can be taken to mean that the person is diseased to these extents, and healthy to the corresponding extents of 0.5 and 0.8. Sadegh-Zadeh (2000) presents a detailed elaboration of

this kind of view using fuzzy set theory.[12] Alternatively, we might picture the vague concept to extend over a spectrum (or a 'space' if the concept is multi-criterial, where the axes represent spectra of different criteria), such that particular cases can be plotted at various locations along the spectrum, to represent their closeness to or distance from uncontroversial cases of disease or non-disease.[13] If the vagueness of the concept 'disease' is understood in such a way, as many-valued, we arrive at an understanding consistent with option 2, implying that there is no definite line to draw.

There are other ways of understanding vagueness apart from the many-valued approach. The 'epistemicist' position regards the exact location of the concept's boundary as unknown; that is, there is a definite boundary there but we do not, and perhaps cannot, know where it is. Alternatively, supervaluationism allows us to consider the disease status of borderline cases to be indeterminate.

Supervaluationism is the view that some claims may be neither true nor false. It tries to solve the problem of borderline cases by admitting that there are cases at the borders of some concepts that do not clearly fall under the concept, but are not clearly excluded by it either. In these cases, a claim that a particular instance does, or does not, fall under the concept cannot be considered to be necessarily either true or false. For example, according to supervaluationism, a statement such as "A man with 500 hairs is bald" cannot be understood straightforwardly in terms of truth or falsity because the concept 'baldness' is vague. There is

[12] Sadegh-Zadeh's (2000) view is much more complex than our simplistic description suggests, incorporating a fuzzy notion of illness as well as of health and disease, and developing a notion of 'patienthood' that is the complement to health. We leave these details aside for the sake of providing a simpler conceptual map of these options, and because they do not affect our argument, but note that Option 2 above is consistent with his (2008) 'prototype' view of disease.

[13] This idea is suggested by Godfrey-Smith's discussion of definitions of biological individuals (2009). We thank Rachael Brown and John Matthewson for pointing us to this work.

a truth-value gap – the claim about baldness is neither true nor false. Similarly, on a supervaluationist view, the statement "X is a disease" is neither true nor false of borderline cases, because the claim would not turn out to be uniformly either true or false on all possible precisifications – reflecting the indeterminacy in our concept (Sorensen 2012, section 5). This way of understanding vagueness supports versions of option 3, insofar as it allows that vagueness does not necessarily imply that there is no boundary, but instead understands the boundary as a borderline zone containing cases about which the truth of claims like "X is a disease" is indeterminate (neither true nor false) pending precisification of the relevant vague concept.

If we accept supervaluationism regarding the vagueness of 'disease', we might continue to regard disease as either present or absent, while also retaining the claim that disease is analysable in terms of dysfunction, by recognising that cases at the borderline between health and disease are simply not categorisable as disease or non-disease, but rather are indeterminate. This kind of view, like the many-valued one, is consistent with allowing stipulative line-drawing about the indeterminate cases for practical decision-making purposes. Thus option 3 retains a categorical distinction between health and disease, while recognising that any particular location for the line will be conventional or stipulative, given that no such line is identifiable by appeal to dysfunction (with or without additional criteria). We could stipulate the location of a line (that is, precisify our concept of 'disease') within the borderline zone as required in relation to various practical or decision-making needs.

This brief discussion of vagueness implies that both options 2 and 3 are potentially plausible. It seems to us that the many-valued approach captures what is going on with conditions where the relevant dysfunction occurs in a step-wise way along a continuous variable, leading

to a relatively wide borderline zone between obvious health and irrefutable disease. Here we are thinking of conditions such as hypertension, hyperthyroidism, osteoporosis or Type 2 diabetes. It makes sense to say that those with 'borderline' readings of blood sugar or bone density have 'a degree' of type 2 diabetes or osteoporosis – and to avoid binary thinking that forces us to define these conditions in relation to a particular threshold. Indeed, using various risk predictors as a proxy, it may be possible to say quite precisely that a person has a specific degree of disease. But in other cases, where the borderline zone is relatively circumscribed, and where it is not possible to track the extent of dysfunction in a linear way across a single definitive parameter, talk of degrees about the borderline cases may not be so apt. This seems to be the case with at least some neoplasms: there is a zone of indeterminacy when the genetic and morphologic characteristics cannot clearly distinguish benign from malignant growth, adjacent to which is a zone where the features are consistent with malignancy on one side, and on the other a zone where no dysplastic cells are present. In these cases, it seems more accurate to say that it is indeterminate whether or not the condition is pathological, rather than to say that cancer is present to some degree, such as 0.2 or 0.5. Indeed, in such cases, whether or not the condition should be counted as a disease often becomes obvious only later, depending on how and whether it progresses, and this is consistent with the medical notion of waiting for a disease to declare itself. At a particular time that information may be ontologically indeterminate (if, for instance, it depends on whether or not some further mutation is acquired that permits the development of a malignant trajectory). And it is certainly, on current medical knowledge, sometimes unknowable. Thus whether or not a particular case is a disease may be neither true nor false. In such cases, supervaluationism seems better to describe what is going on.

In terms of resolving the line-drawing problem, we may wish to draw on option 2 in thinking

about some conditions, and option 3 in thinking about others. Perhaps some sorts of disease admit of degrees, while others are all-or-nothing, but with a borderline zone where, at a point in time, in a particular patient, disease status is indeterminate. Though we do not here argue for any specific approach to conceptual vagueness, a combined approach may be possible, drawing on option 2 in some cases and option 3 in others. We could opt for the many-valued response to vagueness in the general concept of 'disease', and thus allow for degrees of disease, while also saying that some conditions will only ever be present where the degree of disease is '1' (although indeterminate cases at the borderline may occur). Or we might opt for supervaluationism in response to the vagueness of the general concept 'disease', and consider the use of talk of degrees in the case of some diseases to be one way of making stipulative decisions about the presence of disease for practical purposes based on precisification of the concepts involved.

Conclusion

We have argued that dysfunction occurs on a spectrum in a range of diseases including those that seem categorical, such as infections and cancer. As such, dysfunction cannot draw the line between health and disease, creating a tension for accounts that claim both that disease is analysable in terms of dysfunction, and that disease is all-or-nothing. Investigating ways to respond to this tension, we have suggested that the line-drawing problem arises because disease is a vague concept: there are borderline cases that are not clearly disease or health. Different responses to vagueness suggest two plausible options for responding to this tension while retaining dysfunction as a criterion for disease. First, adopting a many-valued approach is consistent with regarding disease to be a matter of degree. Second, adopting a supervaluationist approach allows us to retain the all-or-nothing conception of disease by recognising indeterminate cases at the borderline. Which of these options is preferable

depends on the plausibility of considering diseases as matters of degree, which seems to differ depending on the type of condition. We thus suggest that we might take different approaches when considering different conditions.

Acknowledgements

This research was supported by the Australian Research Council (FT130100346). Our thanks to Thomas Schramme, Yishai Mintzker, HaeJin Jeong, Koji Tanaka, participants at the Defining the Boundaries of Disease conference at Macquarie University in October 2015, and the anonymous reviewers.

References

American Diabetes Association. 2015. Classification and diagnosis of diabetes. *Diabetes Care* 38(Suppl. 1): S8–S16.

American Medical Association House of Delegates Resolution. 2013. Resolution 420 (A-13), http://www.npr.org/documents/2013/jun/ama-resolution-obesity.pdf (accessed December 2015).

Boorse, C. 1975. On the distinction between disease and illness. *Philosophy and Public Affairs* 5: 49-68.

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

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

Boorse, C. 2014. A second rebuttal on health. *Journal of Medicine & Philosophy* 39: 683-724.

Brandal, P., M. R. Teixeira, and S. Heim. 2010. Genotypic and Phenotypic Classification of Cancer: How Should the Impact of the Two Diagnostic Approaches Best be Balanced? *Genes, Chromosomes & Cancer* 49:763–774.

Giesen, L. G., G. Cousins, B. D. Dimitrov, F. A. van de Laar, and T. Fahey. 2010. Predicting acute uncomplicated urinary tract infection in women: a systematic review of the diagnostic accuracy of symptoms and signs. *BMC Family Practice* 11:78, http://www.biomedcentral.com/1471-2296/11/78 (accessed August 2016).

Greaves, M., and C. C. Maley. 2012. Clonal evolution in cancer. *Nature* 481: 306-313.

Hausman, D. M. 2012. Health, naturalism, and functional efficiency. *Philosophy of Science* 79: 519-541.

Hausman, D. M. 2014. Health and functional efficiency. *Journal of Medicine and Philosophy* 39: 634-647.

Hofmann, B. 2016. Obesity as a socially defined disease: philosophical considerations and implications for policy and care. *Health Care Analysis* 24:86–100.

Joint National Committee on Prevention, Detection, Evaluation, and Treatment of High Blood Pressure. 2003. The Seventh Report of the Joint National Committee on Prevention, Detection, Evaluation, and Treatment of High Blood Pressure. *Hypertension* 42:1206.

The published version of this paper is: Rogers W and Walker MJ. The line-drawing problem in disease definition. *Journal of Medicine and Philosophy* 2017; 42 (4): 405-423. DOI: https://doi.org/10.1093/jmp/jhx010.

Lawn, S. D., R. Wood, and R. J. Wilkinson. 2011. Changing concepts of "latent tuberculosis infection" in patients living with HIV infection. *Clinical and Developmental Immunology* 2011: Article 980594.

Lazembnik, Y. 2010. What are the hallmarks of cancer? *Nature Reviews Cancer* 10 (4): 232-3.

Macejko, A.M., and A. J. Schaeffer. 2007. Asymptomatic bacteriuria and symptomatic urinary tract infections during pregnancy. *Urologic Clinics of North America* 34: 35–42.

Moynihan, R. N., G. P. E. Cooke, J. A. Doust, L. Bero, S. Hill, and P. P. Glasziou. 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.

Muir, C. S., and C. Percy. 1991. Classification and coding of neoplasms. In: *Cancer Registration: Principles and Methods* (pp. 64-81), O. M. Jensen, D. M. Parkin, R. McLennan, C. S. Muir, & R. G. Skreet (Eds.). *IARC Scientific Publications* no. 95. Lyon, France: International Agency for Research on Cancer.

Nordenfelt, L. 2007. The concepts of health and illness revisited. *Medicine, Health Care and Philosophy* 10: 5-10.

Pickering, G. W. 1959. The nature of essential hypertension. *Lancet* 1959:1027-1028.

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.

Schmiemann, G., E. Kniehl, K. Gebhardt, M. M. Matejczyk, and E. Hummers-Pradier. 2010. The diagnosis of urinary tract infection. *Deutsches Arzeblatt International* 107(21): 361-367.

Schnarr, J., and F. Smaill. 2008. Asymptomatic bacteriuria and symptomatic urinary tract infections in pregnancy. *European Journal of Clinical Investigation* 38 (S2): 50–57.

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

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

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

Schwartz, P. H. 2014. Reframing the disease debate and defending biostatistical theory. *Journal of Medicine & Philosophy* 39: 572-589.

Sorensen, R. 2012. Vagueness. *The Stanford Encyclopedia of Philosophy*, E. Zalta (Eds.). http://plato.stanford.edu/entries/vagueness/.

Soto, A. M., and C. Sonnenschein. 2004. The somatic mutation theory of cancer: growing problems with the paradigm? *Bioessays* 26(10): 1097-107.

UK Government. 2015. The fit note: A guide for patients and employees. Available at: https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/464395/fit-note-patients-employees-guidance-sept-2015.pdf.

Vickers, A. J., E. Basch, and M. W. Kattan. 2008. Against diagnosis. *Annals of Internal Medicine* 149: 200-203.

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.

Wakefield, J. 1992a. The Concept of Mental Disorder: On the Boundary Between Biological Facts and Social Values. *American Psychologist* 47: 373-388.

Wakefield, J. 1992b. Disorder as harmful dysfunction: A conceptual critique of DSM-III-R's definition of mental disorder. *Psychological Review* 99(2): 232-247.

The published version of this paper is: Rogers W and Walker MJ. The line-drawing problem in disease definition. *Journal of Medicine and Philosophy* 2017; 42 (4): 405-423. DOI: https://doi.org/10.1093/jmp/jhx010.

World Health Organisation. 1948. *International Classification of Diseases, Revision 6*. http://www.wolfbane.com/icd/icd6h.htm, accessed August 2016.