

This is the final author's version of a paper forthcoming in *Philosophy, Psychiatry and Psychology*.

First do no harm? What role should considerations of potential harm play in revising the D.S.M.?

Rachel Cooper

Address: Philosophy; County South; Lancaster University; Lancaster; LA1 4YL; UK

Email: r.v.cooper@lancaster.ac.uk

Tel . +44 1524 594702

Biographical paragraph:

Rachel Cooper is a Senior Lecturer in Philosophy at Lancaster University, U.K. Her publications include *Diagnosing the Diagnostic and Statistical Manual of Mental Disorders* (Karnac, 2013) *Psychiatry and the Philosophy of Science* (Acumen, 2007) and *Classifying Madness* (Springer, 2005).

Abstract

Guidelines for revisions to D.S.M.-5 asked those proposing changes to consider potential harms to patients. This has been controversial. This paper argues that seeking to do no harm is appropriate when constructing a scientific classification scheme. I show that in many cases considerations of harm avoidance can play a role in influencing the design of the D.S.M. without the pursuit of scientific knowledge being compromised. I then turn to cases where compromise is required. Through a consideration of cases where lying is required to save lives I suggest that it is widely accepted that harm-avoidance can trump truth-telling. In extreme circumstances a psychiatric classification might thus set out to mislead in order to prevent harm.

Keywords: classification, *Diagnostic and Statistical Manual of Mental Disorders*, cost-benefit, values, Nassir Ghaemi, lying

First do no harm? What role should considerations of potential harm play in revising the D.S.M.?

Guidelines for revisions to D.S.M.-5 asked those proposing certain types of revision to consider potential harms to patients. Specifically, those proposing new diagnoses were to consider whether "...the harm that arises from the adoption of the proposed diagnosis exceed[s] the benefit that would accrue to affected individuals" (Kendler et al. 2009, 6), and potential for harm was cited as a possible reason for keeping a diagnosis in the appendix rather than promoting it to the main classification (Kendler et al. 2009, 9). The "do no harm" criterion was referred to in many papers arguing for or against proposed revisions to the D.S.M.-5 (e.g. Mataix-Cols et al. 2010, Woods et al. 2010, Shear et al. 2011, Boelen & Prigerson 2012, Huprich 2012, Selby et al. 2012). Types of harm that were commonly considered included stigmatisation and self-stigmatisation, the inappropriate treatment of false-positives, legal and bureaucratic consequences, and the potential medicalization of normality. Conversely, benefits that were commonly expected from proposed revisions included the facilitation of appropriate treatment and other needed services, providing a category useful for future research, and improving clinical communication. For the D.S.M.-5, considerations of harm were only explicitly required from those proposing to introduce new diagnoses, or to move diagnoses from the appendix to the main classification. Still,

presumably, if considerations of harm should play a role in assessing these types of proposed revision, there is no reason why they should not also play a role in evaluating other types of proposed revision (such as proposals for name changes, for lumping or splitting categories, or for altering diagnostic criteria). This paper assesses the legitimacy and application of the “do no harm” criterion. When proposed revisions are considered, should the potential for harm be one factor that is taken into consideration?

The D.S.M.-5 utilization of the “do no harm” criterion has come in for criticism, with Nassir Ghaemi the most vocal opponent (Ghaemi 2010, 2012a, 2013). On his *Psychology Today* blog, Ghaemi pulled no punches. He wrote,

By sacrificing science to ‘pragmatism’, all versions of D.S.M. prevent the profession from identifying diseases, making it harder to help, and ensuring that harm will ensue. (Ghaemi 2013)

In Ghaemi’s view, “we should give primacy to getting at the truth” (2010, 36); a classification should first seek to mirror the natural structure of the domain of mental disorders:

Scientific research in clinical nosology, using classic epidemiological validators (symptoms, course, family history, treatment response, biological markers) should

be the foremost consideration at all times in defining diagnostic definitions. Where there is consensus among scientific experts in a diagnosis, diagnostic considerations should be changed. (Ghaemi 2012a, 51)

This said, Ghaemi allows that considerations of potential harm might play a legitimate role in deliberations about revisions to the D.S.M. in those situations in which the empirical evidence alone does not determine an answer:

Where more or better evidence does not exist, or where significant controversy holds among experts, then practical considerations can be taken into account as a lower-level priority upon which to decide to change or not change diagnostic definitions (Ghaemi, 2012a, 51)

Ghaemi cares about patient wellbeing, but argues that history indicates that in the long run the best way to help patients is to develop a valid classification system that will facilitate research into effective therapies (2010, 2013).

In characterizing the D.S.M.- debate over the relative importance of truth-seeking and harm-avoidance it is important not to over-state the distance between the opposing camps. Everyone involved agrees that scientific research into mental disorders matters, and everyone

also agrees that seeking to minimize harm to patients is important. In a great many cases, both these aims can be pursued together. There is potential disagreement only when it comes to cases in which either pursuing scientific knowledge or avoiding harm will need to be compromised.

The remainder of the paper falls into three sections. In the first, I discuss cases which all should find uncontentious. These are cases in which considerations of harm avoidance may play a role in influencing the design of the D.S.M. without the pursuit of scientific knowledge being compromised. The aim of this section is to show that such cases will be common. As such the potential for harm will frequently need to be considered when possible revisions to the D.S.M. are proposed. The second section considers issues that arise when trying to predict harms: Which of those people who might be affected by a diagnostic change should be considered when possible harms are being judged? Who should judge potential harm? Is it worth trying to judge potential for harm when the future is so uncertain? The third section considers what should be done in those cases where some compromise between pursuing scientific knowledge and harm avoidance may be required. When a choice must be made, which should be prioritized?

1.Cases that should be uncontentious: opportunities for honest do-gooding

Classification requires multiple decisions that cannot be determined by the scientific evidence alone (even in situations in which we have all the scientific evidence we might want). In such cases, striving to minimize harm does not hinder the pursuit of scientific knowledge and should be uncontentious. Consider the following sorts of case.

i. Choices amongst multiple equally evidence-based classifications

In general one should not suppose that only one classification is compatible with the scientific evidence (which might force a choice between a classification that correctly describes reality, or one that is a false representation that will mislead researchers). Reality is complex, and multiple distinct classifications can describe real and important aspects of the world. John Dupré, a philosopher of biology, discusses the various different classifications that biologists use (Dupré 2001). Ecologists favour classifying species in terms of their current behaviour; evolutionary theorists find it more useful to classify on the basis of ancestry. Both ways of classifying result in classifications that truthfully reflect some, but not all, features of reality. Psychopathology is a rich enough domain for multiple different classifications to be consistent with the evidence. Simple cases arise when decisions have to be made whether to lump or split disorders. Frequently the evidence shows that putatively distinct forms of a disorder display both similarities and differences. In such cases, nothing in nature dictates whether the

classification should lump or split. Still, lumping and splitting can have differential effects on patients, and considering the prospects of causing benefit or harm is appropriate; most accept that splitting is particularly justified when linked to differential treatment, for example. Other types of case arise when the groupings that best describe certain similarities simultaneously tend to obscure others. Depending on the interests that one prioritizes different classifications will be most appropriate.

ii. Choices about thresholds

Many disorders are such that there is no “bright line” between mild cases of the disorder and normality; heavy drinking shades into alcoholism, intellectual disability blurs into normal unintelligence, A.D.H.D. merges into poor concentration. In such cases, where should the threshold between normality and disorder be drawn? One reasonable way of fixing the threshold is to choose a point where the likely benefits of diagnosis will exceed the costs. This is the approach that was adopted in determining the number of symptoms that should be required for a diagnosis of A.D.H.D. in D.S.M.-IV, for example, where the cut-off was chosen with the aim of capturing most of those who were impaired to a clinically significant extent (Lahey et al. 1994).

iii. Choices about language and formatting.

What language should be used in describing a disorder? Which diagnostic criteria should be emphasized? How should categories be grouped into chapters? Consider, for example, the decision to rebrand “mental retardation” as “intellectual disability”, and to relegate explicit mention of specific I.Q. scores from the diagnostic criteria to the accompanying text. This change may be hoped to make a difference. The connotations of “mental retardation” and “intellectual disability” are importantly different (Schalock et al. 2007). “Intellectual disability” clearly marks the condition as a disability, and may help those who seek to secure services on a par with physical disabilities. In placing less emphasis on I.Q. scores, the changes that have been made to the D.S.M.-5 may also reduce the frequency with which needy individuals are denied help because their I.Q. test score is slightly higher than the suggested cut-off point. Still, nothing in the science says whether I.Q. scores should be placed in the diagnostic criteria or in the text.

iv. Tougher cases

That considerations of harm might play a role in the above sorts of case is uncontentious. But Ghaemi’s favored example of a case where he thinks that the “do no harm” criterion threatens

to lead to the gerrymandering of categories is tougher. The problem concerns the required duration of hypomania for a diagnosis of bipolar II. D.S.M.-IV required four days. Ghaemi claims that empirical evidence shows that genuine hypomanic episodes (that is, episodes that are similar to longer episodes on important validators) can last for only two days (Ghaemi 2010).

The bipolar II case is doubly complicated at the empirical level: First, there is disagreement as to whether validators suggest that the cut-off for hypomanic episodes should be two or four days (see, for example, Kupfer et al. 2011). Second, there is disagreement as to whether bipolar II is currently over or under-diagnosed. Both over and under-diagnosis could be expected to result in harm. Over-diagnosis leads to unnecessary treatment with side-effect inducing drugs; under-diagnosis means that needy individuals fail to receive medications that could help them. In the run-up to the publication of D.S.M.-5, debates over criteria for bipolar II and the harm criterion became heated, and arguments about the empirical evidence regarding validators, evidence about rates of over and under-diagnosis, and philosophical discussion of the justifiability of the harm criterion came to be mixed together (see, for example, Ghaemi 2010, Pies 2010, Frances, 2010). Here, I am interested in the philosophical question of what should be done when harm-reduction and scientific truth seem to pull in separate directions. To avoid readers being side-tracked by the empirical debates regarding bipolar II we will here switch from talking about bipolar II and take an imaginary case.

Suppose there is a disorder where empirical data regarding validators suggests that diagnostic criteria should be relaxed (e.g. by decreasing a duration requirement, or reducing the numbers of symptoms required) but where there is also robust evidence that the disorder is currently being over-diagnosed and medicated with drugs with dangerous side-effects. What should be done? In such a scenario should the D.S.M. “go with the science”, or fudge the classification in an attempt to do no harm?

At the outset I will dismiss one response to this sort of problem that is sometimes made. It is sometimes claimed that scientists should not have to worry about the foreseeable consequences of their actions. The thought is that scientists should be free to pursue truth and someone else – say policy makers – should worry about how harms can be avoided. In the case of the D.S.M., some might suggest that the committees responsible for the D.S.M. should worry only about constructing a valid classification. If over-treatment is likely to arise this is someone else’s responsibility; in this case maybe medical educators need to do a better job to ensure that diagnostic criteria are applied correctly. The notion that scientists should worry only about truth is considered in detail, and rightly dismissed, by Heather Douglas in her 2003 paper “The moral responsibilities of scientists”. Douglas discusses how fulfilling certain social roles can change an individual’s moral obligations. Usually role responsibilities act so as to confer extra moral obligations. Thus life guards have an extra responsibility to rescue drowning people, and parents have extra responsibilities to care for their children. In some

cases an individual's role can lead to them having fewer moral obligations. Thus a defence lawyer has no duty to report evidence that might implicate her client, although a normal citizen who discovered the same evidence would be expected to report it to the police. Douglas suggests that a role-specific reduction in an individual's moral obligations happens only in cases where a social institution is organised such that some other individual will take up the moral slack. The reasons that the defence lawyer has no duty to report evidence that might implicate her client is that the prosecution team is charged with mounting the prosecution. From this starting point Douglas considers how role responsibilities might affect the moral obligations of scientists. Scientists could only be freed from worrying about the harmful effects of their work if some other group of people could be relied upon to adequately monitor potential harms. However, Douglas argues, in the scientific case such a division of moral labour will seldom be possible, as judging the potential costs and benefits will require scientific expertise possessed only by the scientists concerned. On Douglas's view the reason why it is not acceptable for a rocket scientist to restrict his attention to getting the rockets up, is that there is no similarly qualified person available to worry about where they come down.¹ To apply this reasoning to the D.S.M., the reason why experts on psychiatric classification have a moral duty to consider harms that might result from changes to classification is because they are the only ones who possess the expertise required to judge some of those harms; it takes such experts to know whether some disorder is being over-diagnosed, for example.

Following Douglas we can conclude that scientists have a moral duty to worry about the consequences of their actions. In a situation where there seems to be a conflict between minimizing harm and pursuing scientific knowledge what then should they do? Rather than addressing this question head-on I will take a roundabout route, and first consider how conflicts between publishing scientific findings and minimizing harm have been managed in a quite different scientific field. Switching to consider another case will prove useful, as the practices introduced elsewhere in science might also prove applicable in the case of psychiatry.

The case we will consider concerns a moral dilemma sometimes faced by those scientists who discover new species. Scientific convention dictates that the paper that announces the discovery should include details of the location at which the species was found. Such information is required by other researchers who might themselves want to study the species, and is needed by researchers in biogeography (who, for example, use studies of the range of different species to formulate evolutionary hypotheses and to map the impact of climate change). Others also have an interest in knowing the location of new species; there is a thriving international trade in reptiles and amphibians for the pet trade. When a “charismatic” and collectable species of reptile or amphibian is discovered, commercial collectors can use the scientific paper that announces its discovery as a “treasure map” to locate specimens for sale. In the case of species that only live within a limited area, the effects

of collecting can be catastrophic; a number of species have been pushed to extinction within a few years of their discovery (Stuart et al. 2006).

Those discovering new and charismatic species of reptile and amphibian are thus faced with a dilemma. On the one hand, details of the location at which a species is found are needed for certain types of research. On the other hand, it is now well known that any species where the location is described in the literature is likely to be collected. What is to be done? How might the benefits of potential research be weighed against the potential risks of collecting? In addressing this question we might try to abstract the philosophical essence of the dilemma, and start discussing whether truth or harm-avoidance is the greatest good. This isn't what the biologists did. Rather they considered in greater detail the mechanisms by which both the goods of research and the harms of collecting arise and found a way of both conducting research and also protecting species. Thus, a recent paper announcing the discovery of a new species of tree viper gives only rough details of the location, as "montane forest in Tanzania's Southern Highlands" (Menegon, Davenport and Howell 2011), but provides contact details that researchers may use if their require "additional information on the locality for scientific purposes". The researchers also collected some snakes to set-up a captive breeding population. They plan to give breeding stock away for free to flood the pet-market and destroy incentives for collecting. Given that many reptiles and amphibians can be bred quite easily in captivity,

there is a reasonable chance that this strategy will succeed. If and when it does, the location of the wild snakes might then be publicized with minimal risk.

What can we learn from this case? I think it illustrates that in debates about the relative importance of pursuing scientific truths and avoiding harm talking in general abstract terms (“Is Truth or Harm-Avoidance most important?”) may be neither helpful nor necessary. Depending on the specific details of the situation it may be possible both to minimize the potential for harm and also to enable important scientific work to progress unhindered.

Let’s return to the D.S.M., and our hypothetical case. Suppose there is a disorder where empirical data regarding validators suggests that diagnostic criteria should be relaxed (e.g. by decreasing a duration requirement, or reducing the numbers of symptoms required) but where there is also robust evidence that the disorder is currently being over-diagnosed and medicated with drugs with dangerous side-effects. Following the lead of the biologists the way to deal with this dilemma is to design the D.S.M. so that general clinicians diagnose the disorder only when the stricter criteria are met BUT so that certain types of researcher might yet investigate populations of people who meet less-strict criteria. Careful phrasing can make this achievable, and there will be no need for any dishonesty. This is the approach that the D.S.M.-5 has actually taken, with some distinct diagnostic criteria for research use only included in Section III as “Conditions for Further Study”.

The best thing to do when there seems to be a tension between facilitating science and avoiding harm is to consider in detail what information researchers need and how it is that harms might come about. It may be that a solution can be found which will both minimize the risk of harm and also allow scientific knowledge to be pursued.

2. Issues that arise in seeking to assess potential for harm

The discussion so far should make it clear that there are many cases in which harm-avoidance should be considered when revising the D.S.M. In this section I consider a set of issues that arise when harm is assessed: When assessing possible harms and benefits, whose interests should be taken into account? Who should assess potential harms? Is it reasonable to try and predict harms and benefits when the future is so unpredictable?

2.1 A clarification: Harm to whom?

Those proposing revisions for the D.S.M.-5 were asked to consider potential harms to “affected individuals” (Kendler et al. 2009, 6).² It is somewhat unclear which individuals count as “affected”, and whether, and why, all of them should be considered. In addition to patients, changes to the D.S.M. can also affect other people – family members, those with

shares in pharmaceutical companies, broader society. Should those proposing revisions seek to benefit only patients or also other groups? To worry only about patients seems too narrow a focus. Psychiatry is charged with dealing not only with those who are troubled, but also with a sub-set of those who are troubling. In the case of disorders such as the addictions, paraphilias, and personality disorders, revisions to the D.S.M. should surely consider the impact on broader society, and not simply on those who are diagnosed. On the other hand, it is plausible that not all harms to others should be a concern when the D.S.M. is revised. For example, certain revisions might reduce pharmaceutical company profits, but plausibly this harm can legitimately be ignored. I suggest that revisions to the D.S.M. should consider potential harms only to those to whom psychiatry has a duty of care – most obviously patients, but also broader society.

2.2 Who should judge potential harms?

As already discussed, some of the potential harms associated with altering the D.S.M. can only be judged by experts in psychiatric diagnosis. However, it would be unwise for psychiatric experts to be the only ones who consider whether proposed revisions might do more harm than good. As Philip Kitcher notes in his *Science, Truth and Democracy* (2001) there is a general tendency for researchers to over-emphasize the potential benefits and to under-estimate the

potential harms of their research (Kitcher 2001, 195). In most cases, diagnosed patients are those people who are most likely to be affected by any potential harms (although in some cases, there might be others who would be affected too). As those who will be most directly affected by any diagnostic changes, patients are well-positioned to notice any potential harms or benefits, and should also be consulted.

Occasionally people worry that starting to consider values in science is dangerous. If considering values becomes legitimate, what is to stop “bad” values shaping science? How can we protect against slipping into Soviet-style political psychiatry? Whether value-judgments have some fundamental objective basis is, of course, a deeply contested issue. For the moral realist, those with “bad” values are simply mistaken; the judgment that political dissent should be prevented will be wrong in the same sort of way that an empirical claim might be wrong. Note that even for anti-realists reasoned debate about values can take place against the background of a broadly-agreed framework (Smith 1997). In the case of the possible harms and benefits that might be assessed when the D.S.M. is being revised there is currently very broad consensus. All agree that, for example, it is harmful for someone who is a “false positive” to be treated with drugs that have side effects, and harmful for someone who could be considered merely eccentric to be diagnosed as having a mental disorder. Similarly, all are agreed that where an effective treatment for a harmful condition exists it is a good thing for people who might be helped to be identified. When there are disagreements about the potential

harms and benefits associated with a proposed revision the arguments generally concern empirical details (What is the expected prevalence? What are the long term effects of drug treatment? How stigmatizing is a particular label?). At least in the current context there is no realistic risk that starting to consider values when revising the D.S.M. would open the door to those who would seek to shape the manual to fit unacceptable value-systems; all those involved in the debates broadly agree on what counts as a “harm” or “benefit”.

2.3 A worry about unpredictability

Many are skeptical of attempts to predict the harms and benefits of changing the D.S.M. because they note that the future is unpredictable and that past predictions have often been proved radically wrong (for example, Ghaemi, 2012b, 72). Much depends on anticipated prevalence rates, but these have often been inaccurate. For example, field trials prior to the publication of D.S.M.-IV predicted that changes to the criteria for autism, and the introduction of Asperger’s would make little difference to prevalence rates (Volkmar et al. 1994). In reality, since the D.S.M.-IV there has been a massive increase in autism-related diagnoses (e.g. Pinborough-Zimmerman et al. 2012). Actual prevalence rates prove hard to predict because they depend not only on the contents of the D.S.M., but also on the economic,

cultural, and bureaucratic contexts within which diagnoses are made (e.g. the rise in rates of autism is plausibly related to the provision of better school services).

Further difficulties are caused by the fact that future discoveries are unknown, but can render calculations of costs and benefits obsolete. Consider the potential for future research to reshape debates about the inclusion of attenuated psychosis syndrome (Moran 2009). The concerns of critics are currently framed by the expectation that attenuated psychosis syndrome would be treated with antipsychotics. Consider how the debate might be altered by the outcomes of future research. Preliminary studies suggest that attenuated psychosis syndrome might be effectively treated with fish oils (Amminger et al. 2010). If these findings are confirmed, then the context of the debate would radically shift.

The example of attenuated psychosis syndrome shows how predictions of costs and benefits can be rendered obsolete by future discoveries. Predictions of costs and benefits are also radically context-sensitive in a different sort of way. In weighing up costs and benefits it is necessary to consider factors such as the likelihood that false-positives will lead to harmful drug treatment, whether diagnosis will plausibly lead to helpful early intervention, and the chances that diagnosis will lead to particular groups being stigmatized. The answers to such questions depend on where and how the classification will be used (e.g., in clinical versus community populations, in the U.S. or elsewhere in the world). The impact of a diagnosis of alcoholism, for example, will be quite different for an insured American male, than for a poor

Pakistani female. Context-sensitivity leads to a problem. Before the costs and benefits of revision can be estimated a patient-population has to be assumed. In applying the “do no harm” criterion most of those arguing for revisions for the D.S.M.-5 seemed to have the consequences for clinical U.S. populations in mind. This choice is justifiable if the D.S.M. is thought of as a classification developed primarily for clinical use in the U.S. In so far as the classification is intended for world-wide use, however, the context-sensitivity of many harms and benefits adds to the problems in seeking to estimate whether proposed revisions are likely to do more good than harm.

Given the problems caused by the general difficulty of predicting the future and by context-sensitivity, is it worthwhile seeking to minimise the harms caused by revising the D.S.M.? I suggest that it is. While predicting harms and benefits can be problematic, the impact of at least some types of revision can be predicted with fair accuracy. A particularly straight-forward type of case occurs when there are proposals to medicalise some condition that currently causes no or little harm. Consider, for example, men who like to dress in women’s clothing. If such men are happy they have very little to gain by being diagnosed mentally ill. Quite generally, whenever people’s “symptoms” cause no harm, there is very little to gained from diagnosing them, and in such cases medicalization should be avoided. Costs and benefits can also be fairly well predicted in some other cases. For example, it is a safe bet that lowering thresholds will lead to more people receiving a particular diagnosis,

while tightening exclusion criteria will result in fewer people having a particular diagnosis. Given that our world is not entirely unpredictable, seeking to take into account the anticipated harms and benefits of proposed revisions is worthwhile.

Even when a cost-benefit analysis turns out to be woefully wrong, it is still a good thing that a cost-benefit analysis was attempted. Psychiatry has a duty to care for and respect patients (and, to a lesser extent, broader society), and this implies that organised psychiatry should do its best to minimise harms that might result from its actions. Cost-benefit analyses of proposed revisions can be seen as concrete manifestations of care and respect (Lenman 2000). Framed in such a way, such exercises have a value even in cases where the predictions turn out to be wrong. They demonstrate that those making revisions did their best.

3. Might dishonest do-gooding sometimes be required?

Some will worry that so far I have only dealt with easy cases. It is all very well to argue that harm should be avoided when this is compatible with pursuing scientific knowledge, but what about those cases in which there is a conflict? In cases where harm can only be avoided by obscuring the truth what should be done?

The classic case that is used to illustrate conflicts between harm-avoidance and truth-telling is that where a Nazi comes asking after the Jews hidden in the house. Considering such an extreme case may seem unnecessarily provocative, but is, I think, justified on two counts. First, considering this case is useful because it is already discussed by a body of relevant literature. Second, situations in which psychiatrists have actually been faced with similar dilemmas have arisen; as well as hunting down Jews, the Nazis sought out and killed people with severe mental disorders. Infamously Kant claimed that it was impermissible to lie even to save life. Almost all commentators consider him mistaken, and the dominant view is that lying is justified in the case where a Nazi asks where the Jews are hidden (for discussion see Korsgaard 1986, MacIntyre 1995). As issues regarding conflicts between avoiding harm and truth-telling are too large to consider comprehensively here, I shall simply assume that readers will agree that lying is justified in this case. We will approach the question of whether it might be permissible for a classificatory manual to seek to deceive via considering the similarities and differences between three scenarios in which truth-telling and harm avoidance conflict. Here I do not aim to provide conclusive answers, but merely to point to some features of the scenarios that might be considered relevant. The cases to be considered are as follows:

- A Nazi asks whether any Jews are hidden in the house.

- A Nazi visits a clinic and asks a psychiatrist whether there are any schizophrenic patients.
- A psychiatric classification system is being revised. There is a proposal to add a new sub-type of schizophrenia. It is known that the Nazis will read the manual.

First, let's consider the case where a Nazi visits a clinic and asks a psychiatrist whether there are any schizophrenic patients. This looks very similar to the case where a Nazi asks a householder whether there are any Jews hidden in the house. Indeed, in so far as medics plausibly have special obligations to act in their patients best interests, the psychiatrist has greater reason to lie than the householder.³

The case of the committee that produces a diagnostic manual is slightly different. Suppose that the committee is considering introducing a new subtype of schizophrenia, which it believes to be scientifically justified, but knows that the Nazis will use the classification to identify populations to be killed. Including a restricted as opposed to an expanded category of "schizophrenia" in the manual is not quite analogous to telling the Nazis that there are no Jews in the house. Information given verbally to a Nazi misleads only the Nazi, whereas the information provided by a diagnostic manual will be read by a variety of audiences. While the Nazis might have used the information about a new subtype to kill more people, a research group elsewhere in the world might have used the information to develop more effective

treatments. Still, in certain situations (say, there are no new treatments on the horizon and the Nazis are known to have a killing-program underway) it might be reasonable to think that on balance introducing the new subtype would likely result in very great harm. In general we regard pursuing scientific knowledge as a good, but think that other considerations can be more important. This is why we don't permit experiments that might be informative but that would greatly harm research subjects. I think that preventing many deaths is more important than facilitating psychiatric research. I don't have an argument for this claim, but I take it to be a basic and widely shared moral intuition. Those who agree will think that a manual that will be read by Nazis should avoid expansion of the category of schizophrenia. Note, however, that in practice options short of deceit may well be available and should be used if possible; publication of the new manual might be delayed for some years, the manual might be phrased so as to imply that the subtypes described are only some, but not necessarily all, of those that might exist (as in "subtypes currently recognized are as follows..."), and so on. This brief discussion of the duties of those developing psychiatric classification systems in contexts in which revising the classification might result in extreme harms is admittedly sketchy, but should serve to make it plausible that there may be some circumstances in which avoiding harm will trump the pursuit of scientific truth.

Summary

Guidelines for revisions to D.S.M.-5 asked those proposing certain types of revision to consider potential harms to patients. In this paper I argue that such guidelines are justified, and defend the idea that those revising the D.S.M. should seek to avoid harming patients (and wider society). I first discuss cases in which considerations of harm avoidance may play a role in influencing the design of the D.S.M. without scientific truth being compromised. I argue that such cases are common; they arise when decisions have to be made between distinct classifications that have equal empirical justification (e.g. “lumping” versus “splitting” decisions), and in decisions about thresholds, wording and formatting. Slightly tougher cases, in which there does seem to be a prima facie tension between promoting scientific investigation and avoiding harm can often be dealt with satisfactorily when the details of the case are examined, for example it may be possible for researchers to employ different diagnostic criteria than do clinicians. The second section of the paper considers some issues that arise when potential harms are assessed. I suggested that harms to all those to whom psychiatry has a duty of care should be considered, and that both psychiatric experts and patients should be involved in considering potential harms. Following Lenman (2000) I suggested that cost-benefit analyses of proposed revisions can be seen as concrete manifestations of care and respect. As such, such exercises have a value even in cases where

the predictions turn out to be mistaken, as they demonstrate that those making revisions did their best. The third section considers what should be done in those cases where some compromise between scientific truth and harm avoidance may be required. Through a consideration of cases where lying is required to save lives I suggest that it is widely accepted that harm-avoidance can trump truth-telling. In extreme circumstances a psychiatric classification might thus set out to mislead in order to prevent harm.

Acknowledgements

I am grateful for comments from two reviewers for *Philosophy, Psychiatry & Psychology*. A version of this paper was presented at the British Psychological Society History & Philosophy of Psychology Section Annual Conference. D.S.M.: The History, Theory, and Politics of Diagnosis. University of Surrey, Guildford, March 2013. I am grateful for the comments of those present.

Financial support

This work was supported by a British Academy mid-career research fellowship, MD120060.

Conflict of interest

None.

References:

American Psychiatric Association. 2013. *Diagnostic and statistical manual of mental disorders. (Fifth edition)*. Washington, D.C.: American Psychiatric Association.

Amminger, G., M. Schäfer, K. Papageorgiou, C. Klier, S. Cotton, S. Harrigan, A. Mackinnon, P. McGorry and G. Berger. 2010. Long-chain ω -3 fatty acids for indicated prevention of psychotic disorders. *Archives of General Psychiatry* 67, 146-154.

Boelen, P. and H. Prigerson 2012. Prolonged grief disorder as a new diagnostic category in D.S.M.-5. In *Complicated Grief: Scientific Foundations for Health Care Professionals*, eds. M. Stroebe, H. Schut & J. van den Bout, 85-98, London: Routledge.

Douglas, H. E. 2003. The moral responsibilities of scientists (tensions between autonomy and responsibility). *American Philosophical Quarterly* 40, 59-68.

Dupré, J. 2001. In defence of classification. *Studies in History and Philosophy of Biological and Biomedical Sciences* 32, 203-219.

Frances, A. 2010. “Do no harm” meets “The rule of thirds”, *Bulletin for the Association for the Advancement of Philosophy and Psychiatry* 17, no. 2: 36-38.

Ghaemi, N. 2010. D.S.M.-IV, Hippocrates, and pragmatism: What might have been. *Bulletin for the Association for the Advancement of Philosophy and Psychiatry* 17, no.2: 33-36.

Ghaemi, N. 2012a. Taking disease seriously; beyond “pragmatic” nosology. In *Philosophical issues in psychiatry II: Nosology*, eds. K. Kendler & J. Parnas, 42-53. Oxford: Oxford University Press.

Ghaemi, N. 2012b. Diagnoses as labels. *Philosophical issues in psychiatry II: Nosology*, eds. K. Kendler & J. Parnas, 71-72. Oxford: Oxford University Press.

Ghaemi, N. 2013. D.S.M.-5: If you don't like the effects, look at the causes. *Psychology Today* blog. Posted Jan 9 2013. (<http://www.psychologytoday.com/blog/mood-swings/201301/dsm-5-if-you-dont-the-effects-look-the-causes>). Accessed 26 October 2015.

Huprich, S.K. 2012. Considering the evidence and making the most empirically informed decision about depressive personality disorder in D.S.M.-5. *Personality Disorders: Theory, Research, and Treatment* 3, 470-482.

Kendler, K., D. Kupfer, W. Narrow, K. Phillips and J. Fawcett. 2009. Guidelines for making changes to D.S.M.-V (http://www.dsm5.org/ProgressReports/Documents/Guidelines-for-Making-Changes-to-DSM_1.pdf). Accessed 26 October 2015.

Kitcher, P. 2001. *Science, truth, and democracy*. Oxford: Oxford University Press.

Korsgaard, C. M. 1986. The right to lie: Kant on dealing with evil. *Philosophy & Public Affairs* 15, 325-349.

Kupfer, D., J. Angst, M. Berk, F. Dickerson, S. Frangou, E. Frank, B. Goldstein, A. Harvey, F. Laghrissi-Thode, M. Leboyer, M. Ostacher, E. Sibille, S. Strakowski, T. Suppes, M. Tohen,

R. Yolken, L. Young and C. Zarate. 2011. Advances in bipolar disorder: selected sessions from the 2011 International Conference on Bipolar Disorder. *Annals of the New York Academy of Sciences* 1242, Annals meeting reports 1-25.

Lahey, B., B. Applegate, K. McBurnett, J. Biederman, L. Greenhill, G. Hynd, R. Barkley, J. Newcorn, P. Jensen, J. Richters, B. Garfinkel, L. Kerdyk, P. Frick, T. Ollendick, D. Perez, E. Hart, I. Waidman and D. Shaffer. 1994. D.S.M.-IV field trials for attention deficit hyperactivity disorder in children and adolescents. *The American Journal of Psychiatry* 151,1673-1685.

Lenman, J. 2000. Consequentialism and cluelessness. *Philosophy & Public Affairs* 29, 342–370.

Mataix- Cols, D., R. Frost, A. Pertusa, L. Clark, S. Saxena, J. Leckman, D. Stein, H.

Matsunaga and S. Wilhelm. 2010. Hoarding disorder: A new diagnosis for DSM-V? *Depression and Anxiety* 27, 556-572.

MacIntyre, A. 1995. Truthfulness, lies, and moral philosophers: what can we learn from Mill and Kant?. *The Tanner Lectures on Human Values* 16, 307-361.

Menegon, M., T.R. Davenport and K. Howell. 2011. Description of a new and critically endangered species of *Atheris* (Serpentes: Viperidae) from the southern highlands of Tanzania, with an overview of the country's tree viper fauna. *Zootaxa* 3120, 43-54.

Moran, M. 2009. D.S.M.-V developers weigh adding psychosis risk. *Psychiatric News*. August 21. 44, no. 16: 5.

Pies, R. 2010. The ideal and the real: How does psychiatry escape the D.S.M.-5 "fly-bottle"? *Bulletin for the Association for the Advancement of Philosophy and Psychiatry*. 17 no. 2: 18-23.

Pinborough-Zimmerman, J., A. Bakian, E. Fombonne, D. Bilder, J. Taylor and W. McMahon. 2012. Changes in the administrative prevalence of autism spectrum disorders: Contribution of

special education and health from 2002–2008. *Journal of Autism and Developmental Disorders* 42, 521-530.

Schalock, R., R. Luckasson and K. Shogren. 2007. The renaming of mental retardation: Understanding the change to the term intellectual disability. *Intellectual and developmental disabilities* 45, 116-124.

Selby, E., T. Bender, K. Gordon, M. Nock and T. Joiner. 2012. Non-suicidal self-injury (NSSI) disorder: A preliminary study. *Personality Disorders: Theory, Research, and Treatment* 3, 167-175.

Shear, M. K., N. Simon, M. Wall, S. Zisook, R. Neimeyer, N. Duan, C. Reynolds, B. Lebowitz, S. Sung, A. Ghesquiere, B. Gorscak, P. Clayton, M. Ito, S. Nakajima, T. Konishi, N. Melhem, K. Meert, M. Schiff, M. O'Connor, M. First, J. Sareen, J. Bolton, N. Skritskaya, A.D. Mancini and A. Keshaviah. 2011. Complicated grief and related bereavement issues for DSM-5. *Depression and Anxiety*, 28, 103–117.

Smith, B. H. 1997. *Belief and resistance: Dynamics of contemporary intellectual controversy*.

Harvard: Harvard University Press.

Strous, R. D. 2006. Hitler's psychiatrists: healers and researchers turned executioners and its relevance today. *Harvard review of psychiatry* 14, no. 1: 30-37.

Stuart, B. L., A.G. Rhodin, L.L. Grismer and T. Hansel. 2006. Scientific description can imperil species. *Science*, 312 no.5777: 1137-1137.

Volkmar, F., A. Klin, B. Siegel, P. Szatmari, C. Lord, M. Campbell, B. Freeman, D. Cicchetti, M. Rutter, W. Kline, J. Buijelaar, Y. Hattab, E. Fombonne, J. Fuentes, J. Werry, W. Stone, J. Kerbeshian, Y. Hoshino, J. Bregman, K. Loveland, L. Szymanski and K. Towbin. 1994. Field trial for autistic disorder in DSM-IV. *American Journal of Psychiatry* 151, 1361-1367.

Woods, S., B. Walsh, J. Saksa and T. McGlashan. 2010. The case for including attenuated psychotic symptoms syndrome in D.S.M.-5 as a psychosis risk syndrome. *Schizophrenia Research* 123, 199-207.

Endnotes

¹ To quote the Tom Lehrer song lyrics "Once the rockets are up, who cares where they come down. That's not my department," says Wernher von Braun"

² The full paragraph reads as follows "Is there potential for harm to affected individuals or other groups of persons because of the adoption of this category in DSM-V? Could the harm that arises from the adoption of the proposed diagnosis exceed the benefit that would accrue to affected individuals? In our view, the main concern of DSM-V should be to avoid harm to the affected individuals, but sometimes broader social or forensic issues that involve potential harm to patient or non-patient groups may warrant consideration."

³ During the Third Reich some psychiatrists did lie to the Nazis, but many psychiatrists cooperated with the killing of mentally ill patients (Strous 2006). Many were committed eugenicists and believed that killing their patients would lead to the greater good in the long run. Considering such cases lies outside the scope of this paper; psychiatrists who were motivated to kill their patients "for the greater good" thought of themselves as facing a choice between prioritizing individual or group wellbeing, as opposed to facing a choice between prioritizing harm-avoidance or scientific truth.