

Conceptual Engineering of Medical Concepts

Elisabetta Lalumera, University of Bologna, Department for Life Quality Studies

August 21, 2023

Book chapter, please cite as to appear in Manuel Gustavo Isaac, Steffen Koch & Kevin Scharp (eds.), *New Perspectives on Conceptual Engineering* (Springer, forthcoming).

Abstract

There is a lot of conceptual engineering going on in medical research. I substantiate this claim with two examples, the medical debate about cancer classification and about obesity as a disease. I also argue that the proper targets of conceptual engineering in medical research are experts' conceptions. These are explicitly written down in documents and guidelines, and they bear on research and policies. In the second part of the chapter, I propose an externalist framework in which conceptions have both the explanatory power of psychological concepts and that of semantic concepts. It is likely, however, that human activities and practices distinct from medical research, and regulated by different practices and epistemic rules, call for different targets for conceptual engineering. I conclude by indicating an open agenda of problems for philosophers of medicine interested in conceptual engineering.

Keywords: Cancer, classification, values, medicine, nosology, concepts, conceptions, obesity, experts.

Introduction

There is a lot of conceptual engineering of medical concepts going on, both in medicine and the philosophy of medicine, but it is not usually recognized as such. This chapter aims to substantiate this claim through two case studies, cancer and obesity. I would also like to extract a methodological point about the conceptual engineering of medical concepts, namely that the concepts that can and should be engineered in medicine are, in the vast majority of cases, experts' conceptions. Experts' conceptions are explicitly written down in classifications and guidelines, they are epistemically accessible and public, and they have an impact on policies, research and cure. The considerations I will develop in favour of conceptions in this case here may hopefully shed light on the issues of the feasibility of conceptual engineering, also known as the "implementation challenge", and of pluralism in conceptual engineering. The chapter is structured as follows. In Sections 1 and 2, and Section 3, I illustrate how cancer and obesity, have been redefined in medicine, because of some specific goals. In Section 4, I argue that these are cases of conceptual engineering, and propose a semantic framework including conceptions. the above-mentioned methodological point on the conceptual engineering project from medical cases, before summing up my conclusions in Section 5.

1. Classifying Cancer

In this and the following section, I will illustrate two ways in which the concept of cancer is and has been, modified in medicine and, partially, in the philosophy of medicine, with a view to epistemic but also non-epistemic goals. In section 5 I will explain that we also need the more fine-grained notion of conceptions of cancer to account for ameliorating changes – but for now, the expression “concept of cancer” will do. Let me start by reporting a few facts about what the concept is of¹. The concept of cancer as it is used now in the medical community refers to more than 100 different diseases, characterized by uncontrolled growth of abnormal cells that can take place in many organs of the body. These cells typically pile up in an unstructured mass (tumour), and sometimes also spread to other parts of the body (malignant tumour). A very broad distinction is made among three kinds of cancers: depending on the type of tissue they begin in, these are carcinoma (a cancer that begins in the skin or cover tissues of organs), sarcoma (begins in bone, cartilage, fat, or muscle), and leukaemia (which starts in blood-forming tissue). Cancer may end up in organ failure, long illness, and premature death. Chemotherapy, radiation therapy, immunotherapy, and surgery are now the main interventions available. The choice among such options depends on the site (where the cancer originated), the grade (how the cells look), and the stage (how far it has progressed).

This explains why a good classification of cancers is important on the practical side. Different kinds of cancer respond differently to interventions, and to diagnostic procedures, so to know more about kinds of cancer is to be able to do more about them. How many kinds of cancer there are and how they are differentiated is also crucially relevant to epidemiology, to drug development, to diagnostic and therapeutic solutions based on artificial intelligence, and, last but not least, to institutional and private healthcare policies, and research funding. There is therefore constant and intense scientific work being carried out behind cancer classification, which is regularly updated in the light of new evidence. The official one, issued by the WHO – unofficially called “the Blue Books” – takes primary site as the organizing principle. Now there are 11 volumes in the series, including *Thoracic Tumours* and *Digestive System Tumours*, and so on (Bosman et al. 2010), (Travis et al. 2015)

These classifications are used worldwide, but they are also frequently challenged by experts in different subfields of medicine. An open issue is whether, and if yes then to what extent, classifying cancers according to their molecular and genetic profiles could be feasible, and better than what we have now. Currently, genetic profiles merely integrate traditional classifications, but not uniformly across them – for example, we have robust molecular profiling notably for breast cancers, but not for digestive system cancers or for prostate cancers (about which I will say more below) (Tsang & Tse, 2020; Varghese, 2015). Diagnostic imaging data may also contribute to changing the morphology-based WHO system, especially when it comes to staging and grading, and are partially already doing so (Purysko et al., 2019).

¹ For an updated reference on cancer and cancer treatment, see Russo et al. (2021). An important monograph on cancer from the point of view of the philosophy of science is Plutinsky (2018).

In general, one could be tempted to think that the more we get to know, the more subtypes of cancers, or new cancer subconcepts, will enter the classifications; it is just a matter of time – and the more concept-splitting, the better. In fact, things are not so simple. We saw above on how many levels the decision of lumping together cancer types or, as more frequently happens, splitting a previous unified concept into two or more subtypes may matter. What should count as a new cancer subconcept? The problem is not only epistemic, but also practical, and even ethical. Suppose that a new proposed disease entity, let us say a glioma, a kind of brain tumour, can only be individuated through a very complex and expensive diagnostic procedure (for example, by combining histological description, molecular tests, and analysis with bioinformatics). The new concept would make brain cancer diagnosis more precise. However, most medical communities and practitioners worldwide do not have access to very complex and expensive diagnostic procedures. This technological disadvantage would not prevent them from understanding the new concept and from using it, for example, when discussing post-diagnosis treatment plans. However, neither would they apply the new concept in their everyday diagnostic activities, nor in their research. According to some, the introduction of the new concept would thus increase the already existing healthcare disparities. The general point here is that precision and worldwide shareability are two desiderata for cancer concepts, but they may conflict. Should we require a repertoire of cancer concepts, so to speak, that is “available to all through affordable tools and acquirable knowledge? Or should we embrace complex, detailed (and thus less affordable) genomic tools in this exercise?” (Salto-Tellez & Cree, 2019).

One solution could be to have a plurality of repertoires for cancer concepts, or a plurality of classifications – one for cutting-edge research, one for everyday clinical care and healthcare reimbursements, and so on. To put it simply, prioritizing research results in increasing cancer subconcepts; worrying about healthcare disparities resolves in sticking to common denominator nosologies; and caring more about the practical side of diagnosis and treatment involves conservatism, rather than frequent revision, of cancer concepts. But even within research, a classification that favours the goals of epidemiology is likely to admit a new disease entity individuated by its causes, whereas one that favours the goals of drug development would prefer splitting cancer concepts according to response to treatment (Richiardi et al., 2017). How to choose, and what to prioritize? The main worries with having a plurality of cancer repertoires, however, are that it would then the issue of who coordinates them and how (an epistemic cost), and it would definitely raise the epistemic risk of miscommunication among stakeholders (researchers, clinicians, patients, and institutions) – and communication between them is precisely one of the main aims of having a cancer repertoire at all. Possibly for these reasons, as we have seen, the WHO classification has split down cancer subconcepts in a conservative way. However it might be in the future, the choice here is not just evidence-based, but also philosophical in nature, as cancer concepts are carved out with epistemic goals, but also non-epistemic values in view (Plutinsky, 2018).

This latter consideration points to a question in social epistemology, on the nature of the expertise required in concept splitting. In fact, there is some discussion going on about how diverse, specialty-wise, the

community of medical experts should be when it comes to cancer classification. Should pathologists still have the last word, as they always did? And what about oncologists, molecular biologists, epidemiologists, and imaging specialists? Pathologists know the differences among different kinds of cancer cells, but other expertises may be relevant. The genetic profiling of cancer cells is becoming increasingly possible and useful, so there is a reason why molecular biologists should be involved. Epidemiologists know how frequent a specific tumor is in a population, and how it varies with time, so they may be involved in adding the information that that tumor is rare, or that it ceased to be rare. Oncologists know about the response of cancers to therapies, and they might want to claim, for example, that two genetically or histologically different kinds of cells are classified as the same tumour because they respond exactly to the same treatment. Imaging specialists, such as radiologists and nuclear physicians, are involved in cancer diagnosis and staging, so they know how early different kinds of cancers can be detected, and how they evolve (Salto-Tellez & Cree, 2019). Professional philosophers of science and medicine are, for the moment, not actively engaged in this process – though arguably the discussions going on in medical journals are philosophical². It is plausible that including diverse medical specialties in the revision of cancer nosology, would result in different cancer sub-concepts. Whether the contribution that each one can bring is purely epistemic (different knowledge about cancer), the decision about inclusion is based on social, political, and practical reasons internal to the medical research community. This is another aspect in which the concept of cancer is much less evidence-based than it may seem, *prima facie*.

2. Demarcating Cancer

Differentiating among subtypes of a superordinate concept, however, is just one of the two pillars of an adequate classification system. A superordinate concept is completely specified when all its subtypes are differentiated one from the other, but also – and maybe primarily – when we know how to fix the boundaries of the concept itself. To put it simply, telling apples from pears should come after, or at least go together with, describing what counts as fruit. In the case of cancer, it is crucially important to demarcate what cancer is from what it is not, and from what is not yet, namely from benign tumours and cancer precursor lesions, such as hyperplasia, dysplasia, and *in situ* tumours.

Before illustrating this demarcation problem, it is worth reminding ourselves in this case too that the importance of discriminating cancer from non-cancer and not-yet cancer is epistemic, but also existential and ethical. Whatever the prognosis is, a diagnosis of cancer is often a life-changing event for the person who receives it. A diagnosis of cancer may turn a person into a chronic patient. It may bring fear, stress, anxiety, and deep changes into one's life projects, expectations, or even values. This tends to happen irrespective of

² Readers familiar with the philosophy of psychiatry may have noticed the similarity between issues about classifications of mental disorders and cancer classifications. Differences are also striking, however. To put it simply, the community of specialists in mental conditions is less professionally and methodologically homogeneous than the community of cancer specialists, and psychiatry is much more contested than oncology (or medical imaging, or molecular biology) in its epistemic credentials. Partially as a consequence, philosophers are relatively more involved in discussions about psychiatric concepts, their definition, and subtypes. This, however, is not the place to develop this point further. See Rachel Cooper's chapter in THIS COLLECTION, FORTHCOMING.

how much the person knows about medicine, or even specifically about cancer – along with the scientific literature, such an effect is well described in a quite well-known book, which collects the experiences of a group of cancer specialists and bioethicists who became cancer patients (Dresser, 2012). Thus, on the one hand, it is good to identify cell mutations – and diagnose the disease – soon, because cancer treatment is more effective when it starts in the earliest stages. On the other hand, however, it is bad to identify cell mutations *as cancer* too soon, as overdiagnosis and overmedicalization of non-risky conditions are harmful to the patient for the reason we have given. Subjective harms resulting from overdiagnosis and overmedicalization can turn into clinical risks (Hofmann, 2014). Moreover, when considering healthcare fairness, or the just distribution of resources, they are also ethically wrong, because they divert resources from where they may be needed to where they are futile (Bosman et al. 2010), (Lysdahl & Hofmann, 2020).

Steering a balance between these two desiderata (soon is good, but not too soon is also good) is one of the main philosophical problems posed by the concept of cancer³. Cancer poses many problems, such as how to cure it, how to explain its formation, how to make diagnosis and treatment affordable to everyone, and so on, but this one can be singled out as philosophical in that it involves a decision on the boundaries of the concept, which is underdetermined by evidence, and it should be based on an assessment of epistemic and non-epistemic values – just as we noticed for the case of splitting cancer, in the last section.

Why are the boundaries of the concept of cancer underdetermined by evidence? First, because arguably the concept of cancer is semantically vague, just like the concept of a heap – one grain of rice does not make a heap, one million grains certainly does, but in a range of intermediate cases, there is no (knowable) fact of the matter as to whether the considered amount is a heap or not. This situation is frequent in medicine whenever the pathological condition can be represented by some of the values of a continuous variable. Think of hypertension, for example. Where does hypertension start? And is it a disease? Philosophers of medicine have noticed that in such cases, a line-drawing problem arises (Schwartz, 2007). Usually, the line is drawn by considering the relative risk of death or illness associated with various values of the variable in question. But also, the possibility and efficacy of the intervention (if it works), and even its cost-effectiveness (if it works well enough compared to its cost, evaluated on many levels), matter to the decision on where to draw the line. As the philosopher of medicine Elselijn Kingma (2015) made clear, those are prudential reasons, not facts; if this is so, then the line-drawing problem arising for many medical conditions may be counted as a case of “value intrusion” in a supposed realm of pure facts. Whether values intrude (a term that carries with it a derogatory sense), or whether they are intrinsic to medical research and practice, is a well-known open and big question in itself (Amoretti & Lalumera, 2021).

³ From the point of view of philosophy of science, cancer raises questions including: Is cancer a single disease, despite its heterogeneity in causal pathways and genetic profiles? We have competing general theories of cancer in medicine, is this situation epistemically good? Does mechanistic explanation apply to the case of cancer? Do holistic approaches to explanation fare better? How to define what counts as carcinogenic? For an overview, see (Plutynski 2019).

An objection can be raised at this point, going back to cancer. It points out that the concept of cancer is not vague like the concept of a heap, or the concept of hypertension, because one abnormal cell should count as cancer precisely because of the mechanism of its abnormal growth. The quality, not just the quantity, of cells matters when it comes to cancer, whereas this is not the case for (hyper)tension, or a heap. One response is that we would have no use for such a purely biological concept of cancer. Even if it were an epistemically good concept – so to speak – because it categorizes similar phenomena together, on the grounds of similarity of mechanisms, it would not be a clinically useful concept. A person with two abnormal cells in their lung and a person with metastatic lung cancer do not have much in common from the clinical, or indeed the existential, point of view. The concept of cancer we need in medical research and practice ought to have both a mechanistic or biological ground and clinical utility. This, however, brings us back to the core of the issue: that the concept of cancer ought to have both aspects and which one to prioritize in different contexts is a philosophical decision. The prostate cancer example that I will discuss just below will give evidence to this claim.

A slightly different way to frame the problem of drawing the line between not-yet cancer and cancer is to focus on the distinction between benign and malignant tumours, or proper *cancers*. The difficulty here is not just that of individuating a discrimination point in a continuum, but additionally that of choosing between different criteria for malignancy, namely between histological and clinical criteria. A little medical information is required here. Histological criteria, the traditional ones, are roughly what the pathologist can see. According to these criteria, benign tumours resemble the morphology of the original tissue and never infiltrate surrounding tissue; they might grow expansively but never replace other organs. Moreover, in benign cases there are no tumour cells in blood or lymphatic vessels; in epithelial tissues, any cellular proliferation is limited to the epithelial basal membrane; in mesenchymal tissues cells are typical and the mitotic activity is low or absent; and in lymphoid tissues, the histological architecture is preserved. Malignant tumours from the histological point of view are just the opposite: they infiltrate surrounding tissues or disseminate to distant organs (metastases); they disseminate tumour cells in blood or lymphatic vessels; in epithelial tissues, they invade the basal membrane; in mesenchymal tissues, they display atypical cells; and in lymphoid tissues, there is a change of the normal tissue architecture. Clinical criteria are about what matters to the patient's risk of disability or death. Clinically speaking, benign tumours never infiltrate surrounding tissues, never metastasize, never recur; they might have expansive growth but never replace other organs; and they do not usually affect general body performance. Complementary as they may seem, histological criteria and clinical criteria do not always overlap. It has been shown that they do not, at least in the case of low-grade prostate cancer, encapsulated thyroid tumours originating from follicular cells, low-grade papillary carcinoma of the bladder, intra-mucosal adenocarcinoma of the colon, *in situ* adenocarcinoma of the lung, and clear-cell tubule-papillary carcinoma of the kidney (Esserman et al., 2014).

Let us look in more detail at just the first example, low-grade prostate cancer. With advanced diagnostic tests such as PSA (Prostate-Specific Antigen) screening it is now easy to detect microscopic foci of prostate cancer in men. When PSA screening was introduced around 1990, there was a logarithmic increase in

diagnosed cases, and the early treatment of young patients began; currently, most prostate cancers are diagnosed at the first stage – specifically, at the stage classified as Gleason 6 according to the traditional system. This is not entirely good, however. It was found that overdiagnosis of small prostate cancers had no significant positive effect on mortality. This is because it brought with it a diagnosis of tumours that, in virtue of their characteristics, would never be clinically manifest, that is, are unlikely to cause any harm to the patient – the so-called “indolent tumours”⁴. Intervening on these tumours (with surgery, radiation, or other interventions) counts as an overtreatment or futile treatment. It means that it does more harm than good to the patient. It may cause anxiety, depression, and general loss of well-being, on top of the side effects of the interventions. Overtreatment is now a big problem in medicine, and it also raises considerations of equity and fairness – if a healthcare system provides futile interventions, it diverts resources from useful ones (Esserman et al. 2014) In light of these facts, around 2012–2016, the issue was taken up and discussed by the urology and oncology medical communities in a series of consensus conferences of experts on prostate cancer. One point raised was semantic: should the label “cancer” be dropped altogether for these abnormal cell formations, which are most unlikely to cause harm? It was noted that fear of death resulting from cancer likely plays some role in overtreatment, and removing the scary label could reduce it (Carter et al., 2012). As a result, low-grade (specifically, Gleason 6) formerly prostate cancers were declassified, and it was proposed that they are not called “prostate cancer” anymore, but rather “indolent lesions of epithelial origin” (IDLE) (Epstein et al., 2016).

It is important to note here that the declassified cell formations count as malignant according to the histological criteria – they can extend locally beyond the prostate and invade nerves – though clinically they are not dangerous. In this case, the concept of (prostate) cancer was reshaped according to clinical criteria, which were preferred over histological criteria. Was it the only option available in the light of evidence? Surely not. Another option, favoured by specialists in the US, was (and still is) to change just the grading label – if it is low-grade cancer, let us call it Grade 1 cancer, and not Grade 6 (Epstein, 2022). But when the decision was taken, it was the best option in the light of the evidence, plus the accorded preference for clinical criteria, and considering the ethical principle of non-maleficence (or doing no harm) as the main goal. In these respects, the concept was engineered, or ameliorated, with a view to non-epistemic goals.

The philosopher of medicine Peter Schwarz argues for a similar conceptual change for the case of ductal carcinoma *in situ*, a small breast cancer, which can now be easily diagnosed during routine screening tests. The similarity consists in the conclusion that such lesion cancers should be declassified, just like small prostate cancers, because they are unlikely to develop into metastases, and therefore cause serious harm, whereas the harm of overtreatment is well documented (Schwartz, 2014). The difference with the case of the prostate I have just illustrated is that Schwartz’s argument focuses on the concept of disease, and not on the concept of cancer. Adopting the theory of disease as dysfunction, defended by Christopher Boorse (Boorse,

⁴ Not all small tumors are indolent, and in other cases an early diagnosis and intervention can save a life. What became clear in the last two decades, however, is that not all small tumors should be treated as soon as they are discovered.

1977, 1997), and balancing such conceptual insight with evidence on the prognosis and risk of ductal carcinoma *in situ*, he claims that this kind of malignant cell formation is not a disease, but a condition of increased risk. This does not mean, however, that the condition is clinically irrelevant. In fact, Schwartz holds that:

the judgment of whether a condition is a disease is separate from whether it carries risks and whether there are treatments that can reduce that risk without imposing unacceptable burdens, cost, or new risks. In fact, simply describing [ductal carcinoma *in situ*] as a disease, if done without emphasizing the assessment of risk and benefit in choosing treatment, might interfere with helping individuals determine whether to undertake certain steps for risk reduction. (Schwartz, 2014, 997)

If being a disease does not equate with being clinically relevant, why bother about demarcating disease from non-disease, in this case as in others, then? Because Schwartz thinks – and I agree, without pursuing the point further here – that there are good arguments for viewing medicine as the science that deals with diseases. Some such arguments, in passing, involve considerations about what a just healthcare system should be, and what it should provide freely to citizens in a democratic society, that is, cure and care for diseases and not, for example, medical enhancement. If we want to ensure that limited healthcare resources are equitably distributed, we need a distinction between diseases and other undesirable conditions first.

To quote Schwartz’s conclusion again:

seeing medicine in this way relies on being able to draw lines between normal, healthy conditions and abnormal or pathological ones, and while there are complexities and challenges to such line drawing, it is still central, and essential, to medicine. (Ibid.)

3. Obesity as a Disease

Schwartz’s conclusion above points out one reason why it is important – rather than merely intellectually engaging – to establish whether a condition is a disease or not. Diseases, unlike other undesirable conditions, such as poverty or aesthetically imperfect bodily features, are what medicine should aim at explaining, caring for, and curing when possible⁵. Another closely connected reason is that diseases generally (though not invariably) bring with them a social status, the so-called *sickness* aspect – conceptually distinguishable from the objective, bodily condition (disease proper), and from *illness* (the subjective state of feeling bad). The sickness aspect is a set of rights, permissions, or obligations, both at the personal and at societal level (Amoretti & Lalumera, 2020; Hofmann, 2002). A person with a disease can be granted sick leave from work, receive a pension or other benefits, and can be declared as legally non-responsible, or less culpable, if she commits a crime; she is eligible to receive reimbursements, or free treatments from doctors and hospitals, depending on the organization of the healthcare system of the country she lives in. At the societal level, when

⁵ There is actually a debate going on about how to characterize the goals of medicine, of which this is just one of the possible positions. See, for example, Schramme (2017), Broadbent (2019), and Chapter 2.

a condition is recognized as a disease, research funding can be allocated for studying it, healthcare facilities can be created to cure it, and policies can be implemented to prevent it. This is why fibromyalgia patients in the past, and more recently long-Covid patients, have been advocating the recognition of a disease status – a nosological entry, a proper diagnostic concept – for their condition (Callard and Perego 2021); (Häuser & Fitzcharles, 2018); more examples, such as disphobia and Asperger syndrome, can be found in the recent history of psychiatry (Cooper, 2014). At least in some cases, when a condition becomes a disease – when the concept of that condition becomes a disease concept – there are significant social advantages for affected individuals and groups.

There may also be undesirable consequences of the sickness aspect, though. It is sometimes the case that people with a socially recognized disease status are stigmatized, that is, perceived as different by non-diseased people, and excluded and marginalized based on this. Stigmatization is arguably both morally bad and clinically disadvantageous, as stigmatized people tend to internalize the negative beliefs that society projects onto them, with the result that they give up treatments and renounce healthcare facilities (Goffman, 2009). However, not all diseases are stigmatized – think of the difference between the cancer patient, who tends to be represented as a fighting hero, and the depressed patient, who is definitely not (Corrigan & Bink, 2016; Seale, 2001). Therefore, it is debatable whether stigmatization is merely a consequence of the disease label, or if other variables are in place. We will see this question exemplified in our second main example in this chapter, concerning obesity.

As is well known, obesity has been steadily increasing worldwide in the last 50 years, and now it represents a serious threat to life expectancy in many countries. In 2008, the World Obesity Federation, a medical scientific society based in the United States, published their first position paper about the status of obesity as a disease, including the following explicit definition (italics mine):

[O]besity is a multi-causal chronic disease recognized across the life span resulting from a long-term positive energy balance with the development of excess adiposity that over time leads to structural abnormalities, physiological derangements, and functional impairments. The disease of obesity increases the risk of developing other chronic diseases and is associated with premature mortality. As with other chronic diseases, obesity is distinguished by multiple phenotypes, clinical presentations, and treatment responses. (Allison et al., 2008, p. 1161)

The position has been recently reaffirmed, and the paper in the last version is supplemented with evidence that other scientific societies have endorsed analogous definitions of obesity as a disease in the last decade – specifically 23 of them, located worldwide, and including institutions as diverse as the World Health Organization, the Canadian Medical Association, the European Association for the Study of Obesity, and the Korean Society for the Study of Obesity (Jastreboff et al., 2019, Table 1).

Let me briefly go through the line of argument adopted by the panel of experts of the World Obesity Federation to arrive at their conclusion. First, they assess the possibility of adopting what they called “the

scientific approach”, namely establishing whether obesity is a disease based on objective facts. However, they argue that while there are objective facts of the matter about what obesity is, there are none about what a disease is. In particular, there is a lack of a clear, accepted, and scientifically applicable definition of “disease” that allows one to objectively and empirically determine whether specific conditions are diseases⁶. Then they contemplate the possibility of considering the statements of the majority of relevant experts and societies in the field, which they call “the forensic approach”. The problem with this approach, they conclude, is that it is constitutively insufficient to tell us what is good and right, as different from what is consistent with mainstream opinions. Finally, they decide in favour of what they deem a “utilitarian approach” to the question of whether obesity is a disease, according to which:

conditions that produce adverse health outcomes come to be considered diseases as the result of a social process when it is assessed to be beneficial to the greater good that they be so judged. Such judgments about likely benefit to the greater good are utilitarian judgments that may take empirical input but must also assume certain values. (Ibid.)

Leaving aside the philosophical naivete of the terminology chosen by the World Obesity Federation, the three approaches they consider map onto recognizable philosophical strategies for the task of deciding whether obesity is a disease – and therefore of engineering the concept of obesity, by adding, or not adding, a disease component. The first strategy is what we have seen in the previous section when presenting Schwarz’s proposal about ductal carcinoma *in situ*. Schwartz started from a definition of what a disease is (Boorse’s definition) – a conceptual fact – and then argued that the small breast tumours in question are not diseases, on the ground of such a conceptual fact. This first strategy consists in answering the question of whether the concept of obesity should be engineered into a disease concept, by pointing at the superordinate concept of disease as evidence. The “forensic approach”, on the other hand, is analogous to approaching conceptual engineering as tracking socio-linguistic changes in the use of a term. In other words, this second strategy consists in answering the problem of whether the concept of obesity should be engineered as a disease concept by using sociolinguistic data as evidence (in particular, data about how the concept is used in documents of various institutions and scientific societies). Finally, the “utilitarian approach” is a bad name for the idea that concepts can be ameliorated with a view to non-epistemic goals. It consists in grounding the decision about the concept of disease on considerations of how it would affect the well-being of people with this condition, a fair distribution of research funding, etc. We will return to these issues in the next section.

What are the good effects that would follow from changing the concept of obesity into a disease concept? Here is a list, extracted from the World Obesity Federation paper and statements by the other societies mentioned in the 2019 document. The change would solicit more resources on the treatment and research of obesity; encourage more health professionals to view treating obese patients as a vocation worthy of effort and

⁶ The authors of the position paper consider various definitions of disease defended in the contemporary medical literature or endorsed by scientific associations, such as the American Medical Association. As these mostly recapitulate the debate in the philosophy of medicine about the concept of disease, I refer to (Stegenga 2018) and (Amoretti and Lalumera 2021).

respect; shift public perceptions of obesity away from the view that obesity is a lifestyle choice or aberrant behaviour toward a chronic disease model with individual and population-level determinants; reduce the stigma and discrimination on many obese persons; support policies of regulation for the food industry; facilitate efforts to regulate the marketing of unhealthy foods and drinks to children and adults; and promote structural and environmental changes in workplaces, schools, and communities to prevent obesity and promote physical activity and healthy eating. In sum, most of these good effects have to do with the social aspect of the concept of disease, namely with the rights and obligations that it carries with it (Jastreboff et al., 2019, p. 9).

There is, however, no uniform consensus about whether changing the concept of obesity into a disease concept really amounts to ameliorating it. Or at least, if there is a good consensus within medicine, things are less clear in other scientific and non-scientific communities. Let us look at the disagreements within medicine first. Some scientists claim that obesity cannot be a disease, as its operationalization is flawed. The condition is mainly assessed via the body mass index (BMI), the value obtained from a person's weight in kilograms divided by the square of the height in metres. It is widely recognized that people with the same BMI may be very different, depending on muscle mass, bone density, fat distribution, and other variables; therefore, it is now often claimed that BMI does not individuate phenotypes that are precise enough for research and clinical needs. It is true that scientific societies, such as the World Health Organization, avoid defining obesity via BMI, but in fact, BMI remains the way to measure it in most contexts. For this reason, critics of the proposal that obesity is classified as a disease think that obesity should be assessed in other ways, namely by checking the overall health status of the person. But if this is so, then obesity is not a disease condition per se, but rather it is a disease condition just for those people who are not in good health – which is either trivial or circular (Müller & Geisler, 2017). A possible reply may be that at least some values of BMI are correlated with bad health outcomes beyond a reasonable doubt; however, the methodological point remains open.

A parallel concern about obesity and BMI is voiced by some sociologists and philosophers. It is claimed that “the conflation of measurement with a formal disease label” would encourage stigma rather than reducing it in a society where the idea that bodily appearance is a matter of self-discipline and virtue is deeply rooted (de Vries, 2007) (Gutin, 2021). In fact, the modification of the obesity concept as a disease concept has not yet been tested for its social effects; it is an empirical question whether its implementation will be beneficial or not – as I said above, the association of stigma with the disease label is variable. Moreover, as obesity is correlated with socio-economic status and ethnicity, there is the worry that a disease label would reinforce racism and discrimination. At least in the United States, obesity is prevalent among African Americans and Latinos in general, and in black and Latina/o women and children in particular. The worry is that the obesity diagnosis would strengthen the association of black and brown femininity with pathology, decadence, lack of self-control, and irrationality (Sanders 2019). To this worry, it has been replied that the modified concept would rather help channel resources for health prevention to disadvantaged socio-economic communities – but again, the issue is empirically open.

A further criticism from outside medicine refocuses the problem of the concept of disease involved in the obesity as disease proposal. The philosopher of medicine Bjorn Hofmann (Hofmann, 2016) points out that the conceptual change in question has been decided on pragmatic grounds; that is, it is purely value-based – this is what the position paper of the Obesity Society admits, as we saw above. Starting from this premise, and by testing the obesity-as-disease claim with various philosophical concepts of disease, Hofmann argues that the claim is justified only if a social concept of disease is endorsed. This is the idea that social norms of desirability, advantageousness, and conformity to present-time ideals of efficiency, success, or even morality are conditions for being a disease. As examples of diseases in the social sense, Hofmann proposes homosexuality as it was classified until about 1980⁷, and attention-deficit/hyperactivity disorder (ADHD) now. The upshot of his argument is that there are ethical risks in promoting and applying a social concept of disease, concerning again stigmatization, ableism, and discrimination. If Hofmann is right, then changing the concept of obesity into a disease concept commits the medical community to an undesirable concept of disease:

making obesity a socially defined disease may well turn out to be a wise health policy, improving population health. Especially if we address the social determinants of obesity, such as the food supply and marketing system, disparities in socio-economic status, transportation system, and consumption growth. But to rely too heavily on healthcare may also turn out to be an ineffective and harmful way to solve socially defined challenges. (Hofmann, 2016, p. 97)

Hofmann’s point in the quote above can be put as thus: when a condition is classified as a disease (as in the case of obesity) or de-classified as a disease (as in the case of homosexuality) on value grounds, i.e. to fix social problems connected with the condition, these social problems become healthcare problems, and the risk is that their political and structural dimensions are somehow neglected. To put it very simply, if it’s a disease, then doctors will take care of it – and political institutions can well ignore the issue. This is a point worth considering whenever the concept of disease is at stake, and it adds up to the complexity of the engineering of medical concepts.

4. On Conceptual Engineering

Are you convinced that the obesity as disease concept is better than the obesity as risk condition concept, or is it the other way around? And do you think that histological criteria for cancer should continue to guide classification and demarcation? Whatever the answers may be, my aim here was not to give you the necessary details to arrive at them. My case studies, cancer and obesity, were meant as evidence for the claim that there is conceptual engineering going on in medicine, and in the philosophy of medicine. Conceptual engineering can be defined as “the activity of assessing and developing improvements of our representational devices” (Cappelen & Plunkett, 2020, p. 133). In the examples, we saw how the medical community, or relevant subcommunities within it, proposed to make ameliorative changes to the concepts of cancer and obesity. They faced the question of what concepts of cancer and obesity, respectively, they should have – a normative question (Dutilh-Novaes, 2018; Simion, 2017). In particular, Section 1 about cancer classification

⁷More precisely, the American Psychiatric Association voted in 1973 that “homosexuality ... by itself does not constitute a mental disorder” (Bayer 1987) p. 128). But ego-dystonic homosexuality remained as a disorder in the revised version of DSM-II, published in 1974. The diagnostic entry was definitively removed in DSM-III, issued in 1980. That is why it can be said that homosexuality was “classified” as a disorder until about 1980.

ended up by showing the different and alternative values that may shape the decision to create a new subordinate concept of cancer – a new nosological entity – rather than a fully implemented change. In Cappelen and Plunkett's (2020) terms again, this is a case of non-purely epistemic amelioration, as we saw that new cancer subconcepts are not just subject to the constraint of being more precise, but also more useful, less harmful to patients, and more advantageous for the healthcare sector as a whole. In the section about cancer discrimination, again, we saw a case of conceptual amelioration in view of a practical and ethical goal, namely, to avoid overtreatment as well as clinical and psychological harm to the patient. This was successfully implemented. Finally, the discussion of obesity showed a partially successful conceptual engineering project, again driven by practical and ethical goals.

The discussion on the cancer case and the obesity case could be conducted on different levels. First, they are evidence for a claim that many philosophers of science and medicine have argued for in recent decades, namely that science in general, and medicine in particular, are value-laden in an ineliminable way (Amoretti & Lalumera, 2021; Douglas, 2009; Longino, 1996; Rudner, 1953). Second, they also testify to the role of consensus conferences and expert meetings in the making of medical knowledge – medicine is not just evidence-based, but it is relevantly based on people's agreement about practical and ethical choices (Fanti et al. 2019), (SOLOMON 2016), (Stegenga 2016).

Moreover, they might be taken as evidence against a tendency within the contemporary philosophy of medicine, namely that of discarding as irrelevant any discourse about concepts in medicine – viewed as the old tradition of conceptual analysis, the search for necessary and sufficient conditions for the application of a concept, that is, a definition immune to counterexamples (Fuller, 2018; Lemoine, 2013; Sholl, 2020). I believe that the cases I presented show that the discussion about concepts is alive and kicking in medicine, and it brings about changes. So, it is something philosophers of medicine should pay attention to, and possibly take part in, with discussion and arguments. Authors of the anti-conceptual tendency, however, are right in pointing out that purely descriptive conceptual analysis – with no amelioration in view – is not part of medical research, and this may count as a reason for downgrading it in the philosophers' agenda as well.

In this chapter, however, I intend to illustrate how the cases of conceptual change in medicine I presented bear on the ongoing philosophical debate on conceptual engineering. The rest of this section will then be dedicated to this task. I will start with the questions: What are the representational devices that we can ameliorate and change in the cases I presented? Are these the ideal or proper objects of conceptual engineering in general? This will bring me to illustrate the metasemantic framework for conceptual engineering I favour. To anticipate, I distinguish between psychological concepts, semantic concepts and conceptions, and the latter are supposed to do most of the explanatory work – I believe that conceptual engineering in medical cases like cancer and obesity is best understood as modification of experts' conceptions. I also tentatively propose that other domains of discourse call for other frameworks, some of which do not include conceptions in a central explanatory role. The engineering of the concepts of concept and conception, however, is not my main goal

here, and this section has the more modest one of illustrating one possibility, to be explored further elsewhere. So, let us get down to details.

On the one hand there are psychological concepts. The core idea here is that concepts are constructs that explain people's behaviour in a wide set of cognitive tasks including recognition, categorization, inference, imagination, etc. Psychological concepts are whatever people use by default to identify and categorize individuals or kinds of things – not just beliefs, but also prototypes, exemplars, theories, etc. According to this characterization, a psychological concept of C is a functional kind that can be realized by different psychological entities ((Isaac 2020) (Lalumera 2010) (Machery 2017)). My psychological concept of obesity, for example, may include exemplars of obese persons I have stored in memory, some beliefs about body-mass index thresholds for obesity, as well as the belief that obesity is a contested. Psychological concepts are highly variable among different people (my concept of a kind of thing can be quite different than your concept of the same kind); they can also be diachronically variable within-subject (my concept can become richer or change because of my experiences, testimonial knowledge, and all possible changes in my level of familiarity with the kind in question and expertise).

On the other hand, there are semantic concepts. These are constructs that explain the truth of categorial judgments and the validity of inferences. The semantic concept of category C is what truly applies to all and only things that are C. Its intension, definition, or constitutive principles may diverge, and often do, from what people use fallibly to identify or recognize C. A semantic concept is also (according to many), the meaning of the word (or words in various languages) that the concept expresses. Traditionally, there are three main alternative accounts of how a concept applies to (or stands for, or represents, or is true of) something. Descriptivists hold that the definition of C (the intension), determines the extension of the concept of C, including of all and only things that fit the definition (Evans and Altham 1973). One can relax the point on definitions a bit, and propose that a set of constitutive principles implicitly individuates C (Peacocke 1992). Direct reference theorists deny all this, and insist that what makes something the concept of C is a causal (or more generally, reference-fixing) link with C.

In a somewhat metaphorical, and yet important sense, semantic concepts are not in people's minds, but rather they are grasped by people's minds, sometimes fully and some others partially (and I will say more on this later). For example, I fully grasp the concept of a prime number (I know the definition that generates all and only prime numbers), whereas now I do not fully grasp the concept of, say, the limit of an arithmetical function – though I used to, when I was in high school (Peacocke 1992, 1998). Traditional conceptual analysis can be characterized as the quest for semantic concepts so conceived – usually taking the form of a quest for correct definitions (Laurence and Margolis 1999).

Should conceptual engineering be directed at changing psychological concepts, or semantic concepts? Let me mention here that one could skip the question altogether and remain neutral, proposing formulations like the one I mentioned above, that conceptual engineering is aimed at ameliorating “representational devices”

(Cappelen 2018; see also (Burgess and Plunkett 2013) and (Isaac 2021a) for a critique). In the specific context of medicine and perhaps more generally science, however, I believe that the choice of what to target makes a difference, as I will explain below. So let me dismiss neutrality, and briefly assess the pros and cons of the above-mentioned alternatives.

The main reason in favour of psychological concepts as proper targets is that conceptual engineering aims at real-world changes of people's behaviour, and psychological concepts are what people actually use in the real world when they judge, discriminate, make plans, and so on (Isaac 2021b). One reason against engineering psychological concepts concerns the feasibility of the project – can it really be accomplished? This kind of question in its general version is known as the implementation challenge for conceptual engineering, which concerns semantic concepts, for different reasons⁸. Modifying psychological concepts might turn out to be difficult if feasible at all, considering that they are highly variable intersubjectively and intra-subjectively over time, and they are difficult to get to know about. We would need an empirical study to find out what is the concept of obesity that a group of people employs, and after that, the mechanisms of concept formation and change to intervene on are many and complicated. as feasibility correlates with scope – given a small group of people, in a limited time, it is possible to investigate empirically what their psychological concepts are and to study how to intervene (Jorem 2021).

The second reason against concentrating on psychological concepts is not practical but methodological, and it can be turned into the main reason in favour of semantic concepts as targets of conceptual engineering. With conceptual engineering, we aim at altering the truth value of categorical judgments, such as “this is cancer” or “this is a woman”, or “this is a disease” applied to particular cases, sometimes new cases (Flocke 2021). To do that, we need to change what truly counts as cancer or women or disease, that is, the semantic concept (intension and/or extension), or the meaning of the associated word. The problem with this view has been put in terms of irrelevance: even if we could change the meaning of our target concept, people might go on using their old psychological concepts all along. Precisely because semantic concepts are not in our heads – as externalism requires⁹ -, changing them can be irrelevant to how we reason and act (Koch 2021b). A different but related worry has to do with the metaphysics of semantic concepts; as David Chalmers observes, it's not obvious that you can engineer a new resident of the third realm (Chalmers 2020).

Still, the most ambitious claim that conceptual engineering should be addressed at changing both how people categorize and reason, and what counts as a true categorization judgment containing a certain concept is attractive, even in the simplified version I have just sketched – if we want to change our concepts, we need both the semantics and the psychology¹⁰. What is to do, then? In previous papers (Lalumera 2014) (Lalumera

⁸ The implementation challenge is discussed in a number of recent works, including (Cappelen 2018), (Deutsch 2020), (Deutsch 2021), (Koch 2021a), (Pinder 2021), (Pollock 2021), (Riggs 2019).

⁹ Classical references for meaning externalism of various kinds are (Burge 1979) (Kripke 1972) (Putnam 1975).

¹⁰ See (Koch 2021b) for the same conclusion with detailed arguments. See also (Fischer 2020), and (Isaac 2021b).

2016) I argued that the notion of conception can do some useful theoretical work here (and elsewhere). Conceptions may be the missing link. Though there are diverse notions of conception in the philosophical market, this is the one I favour: conceptions are what people have in mind when they possess a (semantic) concept. They are internalizations of semantic concepts. This characterization of conceptions comes from Christopher Peacocke's work on Fregean concepts (Peacocke 1992) (Peacocke 1998), and its most recent defence is articulated by Sarah Sawyer (Sawyer 2018) (Sawyer 2020) (with a difference I will spell out right below here). Conceptions are part of externalist semantic frameworks, and they are often characterized as sets of beliefs about a category. I prefer the expression "what people have in mind" because it leaves open that nonconceptual and yet semantic content (such as perceptions) be included (this was an essential part of Peacocke's project). Conceptions stand to concepts as many to one. They have normative properties (they can be applied correctly or incorrectly), but they do not fix the extension of a concept. They can be fully explicit, and take the form of a theory, a list of principles, or a definition, but they can also be very sketchy – for example, my conception of inflation may just amount to 'a rise in prices that economists study and that is bad' – and implicit, possibly for example, one's conception of everyday objects like chairs and cups, or game, if Wittgenstein was right (Wittgenstein 2010)¹¹.

The advantage of buying into conceptions is that one can say that different people, or the same person in different times, can have different conceptions, and yet possess and use the same concept. You can have both variation and conceptual sameness, disagreement and topic continuity. You can also make sense of the intuition that everyone's conception of something is wrong, and yet there is something all such conceptions approximate. If you're a causal theorist, in order to account for sameness and continuity you would need to defend the view that having the same concept, for people with different conceptions, comes with a causal link with the relevant category (Sawyer 2020). If you are either a descriptivist or a dual theorist, on the other hand, you would opt for some variant of the idea that different people can understand the constitutive principles of a concept differently -, for example partially, implicitly, or deviantly¹². Yet, insofar as they fulfil the minimal conditions for concept possession, they all refer to the same category, individuated by the definition or by the constitutive principles. This latter is the alternative I favour, but I will not defend it further here, as it does not make any relevant difference to my main claim in this chapter (which is that at least in some relevant cases, conceptual engineering in medicine is engineering of experts' conceptions). Finally, conceptions are in people's minds, so they are not causally inert as semantic concepts are usually considered.

Since psychological concepts are in people's minds, too, I still need to specify the difference between psychological concepts and conceptions in my view, before closing this metasemantic interlude. I take it to be

¹¹ The motivations for including conceptions in one's metasemantic framework are similar to those for adopting a dual-aspect theory of concepts, for which I point again to (Koch 2021b). The difference between conceptions and one of the two aspects of concepts envisaged by the dual theorists are subtle. A quick way to illustrate it could be that conceptions in my view have normative properties, whereas the cognitive psychological parts of dual aspect concepts do not. An extensive discussion of this issue, however, would lead us too far from the main focus of this chapter.

¹² And yet they are entitled to them -this view is articulated in (Scharp 2013).

mainly a difference in explanatory function. Psychological concepts explain default categorization, while conceptions, like semantic concepts, explain truth and correctness of judgments¹³. These can and do come apart - it may well be that what I use for the default categorization of, say, gold is something of the form “shiny and bright yellow metal”, and yet this is not what I take to be the definition or the constitutive principles of gold (however incomplete these might be in my mind). From the epistemic point of view, then, a person’s conception of C can be inferred from the judgments containing C she takes as true. Note that there is room for a distinction between truth and correctness here: someone whose conception of fish includes whales will correctly apply her conception of fish to whales, but she will produce a false judgment. So, unlike psychological concepts, conceptions can be applied correctly or incorrectly – they have normative properties. Unlike semantic concepts, however, conceptions are part of people’s mental furniture and thereby they matter to what single agents do and think – so that the irrelevance objection to the claim that they should be the targets or conceptual engineering loses at least part of its grip.

I finally arrived at proposing that what goes on in medical research in cases like those I discussed – cancer and obesity – is conceptual engineering targeted at expert conceptions, or more precisely, at the conceptions that experts agree on. Here are the explanatory virtues of experts’ conceptions in this context. They are always explicit; they take the form of outright definitions or constitutive principles. These are proposed and discussed in consensus conferences and published in articles and guidelines. Therefore, they are not ineffable, but rather epistemically easy to reach; and they are public. These two characteristics make medical experts’ conceptions something it is possible to act on. Also, they are implemented in actions – in research, care, cure, and healthcare policies. So, they are something it is worth acting on (again, no irrelevance objection lurks in here). It can still be objected that on a strictly externalist view (not a socio-externalist one) conceptions per se do not determine meaning, there is always a logical gap between use and meaning. So when experts’ modify their conception, it is not the case that the extension of cancer or obesity ipso facto changes. Moreover, after all, experts can be wrong. I have a concessive reply to this worry. In the particular case of science and medicine, the correctness of experts’ judgments (“Covid-19 is caused by a virus”, “Pluto is not a planet”) is generally considered to be a proxy for truth. Experts’ conceptions do not fix what’s true, but still – on this strongly externalist view – they fix a socially recognized semantic standard of correctness. A change in publicly recognized standards of correctness can be good enough for the practical purposes we aim at – this is how we can obtain semantic change, albeit in this proxy form, in an externalist framework¹⁴. Notice again that this proximity of correctness to truth is a peculiar feature we ascribe to scientific judgments, and it need not generalize to other less hierarchically structured knowledge communities.

Again, a few words are in place to recap why neither psychological concepts nor semantic concepts would be the ideal target of conceptual engineering here. It is surely true that how people think about their conditions (psychological concepts) has an impact on what counts as a disease, on what counts as urgent

¹³ This might apply to the cognitive contents of dual aspect theories of concepts as well, see footnote 11 above.

¹⁴ For a related response to the externalist problem, see (Chalmers 2020).

research, and on healthcare policies – to cite a few aspects only. Consider the case of homosexual activism that led to the (slow but definitive) de-pathologization of homosexuality in the late Seventies of last century. People’s ways of thinking did make a change in that situation. However – I submit - they made a change by altering experts’ conceptions, that is, what figures in nosologies, guidelines, research projects, and policies. It was only when a new conception of disorder and homosexuality was adopted by the experts, that real-world effects became regular and stable on the life of homosexual people. The hierarchical distribution of knowledge and power in medicine explains this fact. This by no means denies the possibility that experts’ decisions, and therefore experts’ conceptions, be arrived at in participatory and inclusive ways, so that what people think is given a recognized epistemic and political role in the engineering process. In fact, the tendency now is to advocate for more inclusive expert panels in medicine, and philosophers of science are very active in the project (Bueter 2019) (Tekin 2022). My point here is just that changing people’s psychological concepts alone – granted that this is possible, for the sake of the argument – would not be sufficient in the context of medical research, for structural reasons.

So why not just say that word meanings, or semantic concepts, are what changes when oncologists propose radically different ways of categorizing tumours, or the Obesity Society announces that obese people are diseased? As I argued above, first because if one takes the ontology of semantic concepts seriously, they are abstract entities, and need to be internalized to do some explanatory work in thought and action; second, because of the externalist intuition that expert’s conceptions - as said above - even the more recent and refined ones, about complex biological conditions such as diseases may just be false, and still count as our best effort at grasping what is true of them.

As conceptions are not currently very popular in the philosophical debate, let me point out in conclusion that my proposal is not as revisionary as it may seem, because of its limited scope. I suspect that the distinction between concepts and conceptions is particularly useful when we want to understand and change the dynamics of scientific discourse, for two reasons. First, because the externalist intuition, that concepts are of kinds of things out there that different people relate with and aim at knowing better, with the possibility of massive error, is arguably more compelling when we talk about cancers or mushrooms, than when we talk about, say, cultural appropriation or privacy. Second, because the concept-conception framework works well at explaining different levels of expertise on the same topic, and science is typically organized hierarchically when it comes to knowledge. *Prima facie*, however, it is not plausible that the rules of the game of medical research and healthcare management are the same as those of other human activities.

The thought here is that practices, knowledge hierarchies, relations of epistemic trust and deference, and metasemantic rules for change that hold within medicine (or maybe, other specialized sciences) need not be generalizable to other communities. I would rather speculate that in order to ameliorate teenagers’ concept of obesity, so that it becomes less value-laden and discriminatory, a top-down operation on explicit definitions or constitutive principles would be pointless – where are such definitions supposed to be, in social media posts?

Who's the authority? How are the boundaries of the community itself fixed? If this suspicion is well-founded, then it might be that different metasemantic frameworks can be chosen by conceptual engineers, for different domains and projects – for politics, morality, scientific contexts, or philosophy. Possibly in some of those, conceptions would be explanatorily idle. In fact, in the previous paragraphs I have characterized semantic concepts, psychological concepts and conceptions as constructs, meaning that what counts primarily is their explanatory function within a theoretical project. Taking my suspicion at face value, moreover, means opting for a radical pluralism about the objects of conceptual engineering (Belleri 2021) (Isaac et al. 2022) (Sawyer 2021). My pluralistic attitude extends to the prospects of success of conceptual engineering projects. Speculatively, epistemically hierarchical communities with tight links between members, and explicit statements and discussions on basic conceptual issues, such as that of medical research, look like an ideal field for conceptual engineering. Experts' conceptions are, so to speak, open to view, and so are epistemic rules of argumentation and deference. I am not in a position to project this optimism to other domains, though. My considerations about pluralism, however, stand in need of a more thorough discussion, involving cases of non-medical concepts.

5. Conclusion

In this chapter, I illustrated two cases of ameliorative conceptual change in medicine involving cancer and obesity. In both cases, the amelioration project is approached with practical and ethical goals in view – at the bottom line, with the goal of doing less harm to patients, and having a better quality of healthcare. In the case of cancer, I showed how classifying, and demarcating involves choosing between different evidence-based criteria in a way that cannot be itself evidence-based, but rather value-based. In the case of obesity, I described the proposal of classifying obesity as a disease, aimed at bringing advantages to obese people in terms of more research on their condition, better care, and less stigma. My point in both cases was not to adjudicate whether the proposals are really conducive to amelioration or not. Rather, I wanted to show that medical research and healthcare organization are active fields of conceptual engineering.

In the remaining part of the chapter, I proposed that the most adequate metasemantic framework in the case of conceptual changes in medical research includes the notion of conception, and conceptual engineering in such area is better described as directed to expert conceptions. So, all my previous talk of concepts in the chapter ought to be replaced with conception talk. Expert's conceptions are incessantly rediscussed and reshaped, not only in the light of new evidence, but also and maybe primarily in the light of practical and ethical preferences and values. I argued that experts' conceptions are the preferred and a relatively easy target for conceptual engineering in the medical domain, in that they are public and explicit – therefore addressing the implementation problem for conceptual engineering. Medical and healthcare research, however, is a peculiarly structured activity, and different human practices and communities may call for different strategies for ameliorating our representational devices. This latter consideration opens up the possibility of a plurality of frameworks for conceptual engineering (Isaac et al. 2022).

This chapter is intended as the first step in a not-much-explored territory for philosophers of medicine interested in conceptual engineering. The metasemantic framework I proposed can be objected to, but I hope I pointed to a recognizable area of problems and work to be done. First, we need to investigate the details of how medical concepts or conceptions are formed and changed in various cases (a task which is close to the social epistemology of science and medicine). Here, a reflection on the non-epistemic basis of medical knowledge is in place: what is the role of values in such decisions? Which and whose values are involved? Starting from such investigations, then, there is room for a normative project, that of individuating the principles according to which such changes should be accomplished in general – including norms and criteria about inclusion and public participation¹⁵.

References

- Allison, D. B., Downey, M., Atkinson, R. L., Billington, C. J., Bray, G. A., Eckel, R. H., Finkelstein, E. A., Jensen, M. D., & Tremblay, A. (2008). Obesity as a disease: A white paper on evidence and arguments commissioned by the council of the obesity society. *Obesity*, 16(6), 1161–1177.
- Amoretti, M. C., & Lalumera, E. (2020). The concept of disease in the time of COVID-19. *Theoretical Medicine and Bioethics*, 41(5), 203–221.
- Amoretti, M. C., & Lalumera, E. (2021). Wherein is the concept of disease normative? From weak normativity to value-conscious naturalism. *Medicine, Health Care and Philosophy* online first, <https://doi.org/10.1007/s11019-021-10048-x>
- Belleri, D. (2021). On pluralism and conceptual engineering: Introduction and overview. *Inquiry*, 0(0), 1–19.
- Boorse, C. (1977). Health as a Theoretical Concept. *Philosophy of Science*, 44(4), 542–573.
- Boorse, C. (1997). A Rebuttal on Health. In J. M. Humber & R. F. Almeder (eds), *What Is Disease?* (pp. 1–134). Humana Press.
- Broadbent, A. (2019). *Philosophy of medicine*. Oxford University Press.
- Carter, H. B., Partin, A. W., Walsh, P. C., Trock, B. J., Veltri, R. W., Nelson, W. G., Coffey, D. S., Singer, E. A., & Epstein, J. I. (2012). Gleason Score 6 Adenocarcinoma: Should It Be Labeled As Cancer? *Journal of Clinical Oncology*, 30(35), 4294–4296
- Chalmers, D. J. (2020). What is conceptual engineering and what should it be? *Inquiry*, 1–18.
- Cooper, R. (2014). *Diagnosing the diagnostic and statistical manual of mental disorders* (pagg. xvi, 79). Karnac Books.
- Corrigan, P. W., & Bink, A. B. (2016). The Stigma of Mental Illness. In *Encyclopedia of Mental Health* (pagg. 230–234). Elsevier. <https://doi.org/10.1016/B978-0-12-397045-9.00170-1>
- de Vries, J. (2007). The obesity epidemic: Medical and ethical considerations. *Science and Engineering Ethics*, 13(1), 55–67.
- Douglas, H. (2009). *Science, Policy, and the Value-Free Ideal*. University of Pittsburgh Press.
- Dresser, R. (2012). *Malignant: Medical Ethicists Confront Cancer*. Oxford University Press.
- Dutilh-Novaes, C. (2018). Carnapian explication and ameliorative analysis: A systematic comparison. *Synthese*, 1–24.
- Epstein, J. I., Egevad, L., Amin, M. B., Delahunt, B., Srigley, J. R., & Humphrey, P. A. (2016). The 2014 International Society of Urological Pathology (ISUP) consensus conference on Gleason grading of prostatic carcinoma. *The American journal of surgical pathology*, 40(2), 244–252.

¹⁵ My thanks are due to the Editors of this collection, for very useful comments to the first version of this Chapter, who made me rethink carefully about my metasemantic framework, and to Stefano Fanti for invaluable advice on all things medical.

- Esserman, L. J., Thompson, I. M., Reid, B., Nelson, P., Ransohoff, D. F., Welch, H. G., Hwang, S., Berry, D. A., Kinzler, K. W., Black, W. C., Bissell, M., Parnes, H., & Srivastava, S. (2014). Addressing overdiagnosis and overtreatment in cancer: A prescription for change. *The Lancet Oncology*, 15(6), e234–e242
- Fanti, S., Oyen, W., & Lalumera, E. (2019). Consensus Procedures in Oncological Imaging: The Case of Prostate Cancer. *Cancers*, 11(11), 1788.
- Fuller, J. (2018). What are chronic diseases? *Synthese*, 195(7), 3197–3220.
- Goffman, E. (2009). *Stigma: Notes on the Management of Spoiled Identity*. Simon and Schuster.
- Haslanger, S. (2020). How Not to Change the Subject. In Teresa Marques and Åsa Wikforss (eds.) *Shifting Concepts*. Oxford University Press, 236-255.
- Häuser, W., & Fitzcharles, M.-A. (2018). Facts and myths pertaining to fibromyalgia. *Dialogues in Clinical Neuroscience*, 20(1), 53–62.
- Hofmann, B. (2002). On the Triad Disease, Illness and Sickness. *The Journal of Medicine and Philosophy: A Forum for Bioethics and Philosophy of Medicine*, 27(6), 651–673.
- Hofmann, B. (2014). Diagnosing overdiagnosis: Conceptual challenges and suggested solutions. *European Journal of Epidemiology*, 29(9), 599–604.
- Hofmann, B. (2016). Obesity as a Socially Defined Disease: Philosophical Considerations and Implications for Policy and Care. *Health Care Analysis*, 24(1), 86–100.
- International Agency for Research on Cancer (1965-2021), <https://tumourclassification.iarc.who.int/welcome/> accessed on 28 December 2021
- Isaac, M. G. (2020). How to conceptually engineer conceptual engineering? *Inquiry*, 1–24.
- Jastreboff, A. M., Kotz, C. M., Kahan, S., Kelly, A. S., & Heymsfield, S. B. (2019). Obesity as a Disease: The Obesity Society 2018 Position Statement. *Obesity*, 27(1), 7–9.
- Lalumera, E. (2014). On the explanatory value of the concept-conception distinction. *Rivista Italiana di Filosofia del Linguaggio* 8(3), 73-81.
- Lemoine, M. (2013). Defining disease beyond conceptual analysis: An analysis of conceptual analysis in philosophy of medicine. *Theoretical Medicine and Bioethics*, 34(4), 309–325.
- Longino, H. E. (1996). Cognitive and Non-Cognitive Values in Science: Rethinking the Dichotomy. In L. H. Nelson & J. Nelson (eds), *Feminism, Science, and the Philosophy of Science* (39–58). Springer Netherlands.
- Lysdahl, K. B., & Hofmann, B. (2020). Overutilization of Imaging Tests and Healthcare Fairness. In Lalumera, E. and Fanti, S. *Philosophy of Advanced Medical Imaging*, Springer, 99-111
- Machery, E. (2009). *Doing without Concepts*. Oxford: Oxford University Press.
- Machery, E. (2017). *Philosophy within its Proper Bounds*. Oxford: Oxford University Press
- Plutinsky, A. (2018) *Explaining cancer: Finding order in disorder*. Oxford University Press.
- Purysko, A. S., Magi-Galluzzi, C., Mian, O. Y., Sittenfeld, S., Davicioni, E., du Plessis, M., Buerki, C., Bullen, J., Li, L., Madabhushi, A., Stephenson, A., & Klein, E. A. (2019). Correlation between MRI phenotypes and a genomic classifier of prostate cancer: Preliminary findings. *European Radiology*, 29(9), 4861–4870.
- Richiardi, L., Barone-Adesi, F., & Pearce, N. (2017). Cancer subtypes in aetiological research. *European Journal of Epidemiology*, 32(5), 353–361.
- Rudner, R. (1953). The Scientist Qua Scientist Makes Value Judgments. *Philosophy of Science*, 20(1), 1–6.
- Russo, A. et al. (2021) *Practical Medical Oncology Textbook*. Springer Nature.
- Salto-Tellez, M., & Cree, I. A. (2019). Cancer taxonomy: Pathology beyond pathology. *European Journal of Cancer*, 115, 57–60.
- Sawyer, S. (2020). Truth and objectivity in conceptual engineering. *Inquiry*, 63(9–10), 1001–1022.
- Sawyer, S. (2021). Concept pluralism in conceptual engineering. *Inquiry*, 0(0), 1–26.
- Schramme, T. (2017). Goals of medicine. In Schramme, T., & Edwards, S. (Eds.). (2017). *Handbook of the Philosophy of Medicine*, Springer Netherlands, 121-128.
- Schwartz, P. H. (2007). Defining Dysfunction: Natural Selection, Design, and Drawing a Line. *Philosophy of Science*, 74(3), 364–385
- Schwartz, P. H. (2014). Small Tumors as Risk Factors not Disease. *Philosophy of Science*, 81(5), 986–998. <https://doi.org/10.1086/678280>
- Seale, C. (2001). Sporting cancer: Struggle language in news reports of people with cancer. *Sociology of Health & Illness*, 23(3), 308–329.
- Sholl, J. (2020). Health in Philosophy: Definitions Abound but a Theory Awaits. In J. Sholl & S. I. S. Rattan (A c. Di), *Explaining Health Across the Sciences*, 79–95. Springer International Publishing.
- Simion, Mona (2017). The ‘should’ in conceptual engineering. *Inquiry*, 1–15.

- Stegenga, J. (2016). Three Criteria for Consensus Conferences. *Foundations of Science*, 21(1), 35–49.
- Tsang, J. Y. S., & Tse, G. M. (2020). Molecular Classification of Breast Cancer. *Advances in Anatomic Pathology*, 27(1), 27–35.
- Varghese, A. (2015). Chemotherapy for Stage II Colon Cancer. *Clinics in Colon and Rectal Surgery*, 28(4), 256–261.