On the uses and abuses of biomarkers in clinical reasoning

Benjamin Chin-Yee

Abstract

Biomarkers are central to the practice of precision oncology, which looks to novel biomarkers to 'personalize' cancer care. Philosophers have highlighted epistemic issues surrounding biomarkers but a general account of their role in clinical reasoning is lacking. This article examines biomarker use in clinical reasoning through the lens of abstraction. I propose *clinical abstraction* as a descriptive and normative account of reasoning with biomarkers that overcomes epistemic and ethical problems raised in the literature.

1 Introduction

Consider some sentences you might hear on any given day in a busy cancer clinic:

> *This is a case of myelodysplastic syndrome with 5q deletion.*
>
> *This is a case of acute myeloid leukemia with NMP1 mutation.*
>
> *This is a case of chronic lymphocytic leukemia with 17p deletion.*

Each prepositional phrase is rich with information. Each carries special diagnostic, prognostic, and therapeutic significance. Their prepositional objects—*5q deletion*, *NPM1 mutation*, *17p deletion*—are biomarkers. Genomic biomarkers, specifically.

Biomarkers are biological data that serve as objective indicators of the state of a patient. For genomic biomarkers these data are genomic—measures of genes and their effects. Genomic biomarkers play a key role in precision oncology which looks to novel biomarkers to 'personalize' cancer treatment.[1] Together, these sentences signal the arrival of genomic biomarkers in routine clinical practice where they are having a profound impact—reshaping cancer diagnosis, prognosis, and treatment.

Recently, scholars in philosophy of science and science and technology studies have examined the epistemic, ethical, and social dimensions of biomarkers in oncology research and practice, including how they shape patient classification, clinical decision-making, and experiences of uncertainty (e.g., Hey 2015; Blanchard 2016; Timmermans et al. 2017; Kerr et al. 2019; Dam et al. 2022). This work has also highlighted challenges and potential harms posed by biomarkers in clinical settings (e.g., Hey and Barsanti-Innes 2016; Plutynski 2022; Reynolds 2020). Clarifying these concerns, however, requires a better understanding of the function of biomarkers in clinical reasoning. How should we interpret the above sentences in practice? What does it mean to call a patient *a case of chronic lymphocytic leukaemia with 17p deletion*? To whom should we apply this label? What special inferences and actions, if any, does it warrant?

[1] See, for discussion, Plutynski (2022) and Chin-Yee (2024).

Answering such questions calls for a more general account of the role biomarkers play in medical practice. Absent such an account, epistemic and ethical concerns surrounding genomic biomarkers in precision oncology will remain difficult to resolve.

In this article, I develop an approach to reasoning with biomarkers called *clinical abstraction*. My central claim is that we should interpret the above sentences as *abstractions* which serve as tools for reasoning and action in particular cases. Prima facie, this interpretation might seem trivial. Insofar as these sentences highlight certain properties and omit others they indeed, uncontroversially, qualify as abstractions. But the implications of this view are far from trivial. As I show in this article, this interpretation is not only key to addressing specific epistemic challenges surrounding genomic biomarkers in precision oncology but also offers important insights for clinical reasoning more generally.

The article is divided into two sections. §2 begins with three illustrative cases before introducing definitions of biomarkers and their role in clinical practice. I address concerns over epistemic harms posed by genomic biomarkers, arguing that these concerns are better understood under a more general account which sees reasoning with biomarkers through the lens of abstraction. §3 develops the approach to reasoning with biomarkers that I call *clinical abstraction*. I contrast this approach with scientific abstraction and highlight its advantages in medical practice. I conclude by applying *clinical abstraction* to one of the cases to show how it offers both a descriptive and normative account of clinical reasoning with biomarkers.

This article focuses primarily on how to reason with genomic biomarkers in the practice of precision oncology where I argue that thinking in terms of clinical abstraction is particularly helpful to address emerging epistemic challenges. I believe my account offers broader lessons, beyond biomarkers, on how to reason with abstractions in medicine, though full discussion is beyond the scope of this article.

2 Biomarkers in clinical practice

2.1 Cases

Dr. Roy, a resident, is working with Dr. Cong, the consultant in the hematology clinic. They see the following three cases.[2]

> THROMBOSIS. The first patient, Mrs. de Vries, has deep vein thrombosis, a blood clot in her leg. A quick study, Dr. Roy knows the state-of-the-art genomic biomarkers that predict risk of blood clots and astutely applies these tests. He eagerly presents the case to Dr. Cong: '*This is a case of deep vein thrombosis in a Factor V Leiden heterozygote*. Based on this mutation, the patient's risk of recurrent blood clots is around 30 percent over five years. We should recommend lifelong blood thinners to decrease this risk'. Dr.

[2] The clinical cases described in this article are anonymized composite examples drawn from the author's clinical experience. The names and dialogues are fictive, and any potentially identifying details have been altered or removed. As such, this work was exempt from Research Ethics Board review as per institutional policies.

Cong, shaking her head, corrects him—'You mean to say, this is a case of *provoked* deep vein thrombosis *following hip replacement*. Because the clot was provoked by major surgery, Mrs. de Vries's risk of recurrence is probably closer to 3 percent. We should recommend only three months of blood thinners to start'.

GENETIC TESTING. The second patient, Ms. Ali, was referred for abnormalities on genetic testing. The resident, Dr. Roy, is again familiar with the state-of-the-art genomic biomarkers and their disease associations. He presents the case to the consultant, Dr. Cong: '*This is a patient with a mutation in the TP53 gene* which causes several cancers, including leukemia. We should inform the patient of this abnormality and her risk of developing leukemia'. Dr. Cong, frowning, interjects—'You mean to say, this is a patient *with a variant of uncertain significance* in the *TP53* gene. Many such variants turn out to be benign. We should perform more tests before counselling Ms. Ali on her risk'.

LEUKEMIA. The third patient, Mr. Lund, was recently diagnosed with chronic lymphocytic leukemia, a common type of blood cancer. Ever keen to impress, Dr. Roy again applies the state-of-the art genomic biomarkers. He confidently presents the case: '*This is a case of chronic lymphocytic leukemia with 17p deletion*. This genomic alteration causes rapid disease progression. This patient has a poor prognosis and would benefit from upfront treatment with a targeted drug'. To the resident's dismay, Dr. Cong again interrupts—'You mean to say, this is a case of *asymptomatic, early-stage* chronic lymphocytic leukemia with 17p deletion. The prognostic significance of the 17p deletion in this context is unclear. We should gather more information before deciding if Mr. Lund would benefit from treatment'.

* * *

Where did the resident Dr. Roy go wrong in these cases? There are different ways to describe the resident's errors. One way is in how he categorized the patients. In THROMBOSIS he used the reference class 'cases of deep vein thrombosis in *Factor V Leiden heterozygotes*' rather than 'cases of deep vein thrombosis *following major surgery*' to estimate Mrs. de Vries's risk of recurrent blood clots. The consultant, Dr. Cong, recognized that the latter category provided the best estimate of the patient's risk of recurrence and treated accordingly.

A second way to describe the resident's errors is in how he assumed homogeneity of a given reference class. In LEUKEMIA, Dr. Roy assumed that poor prognosis is a uniform feature of 'cases of chronic lymphocytic leukemia *with 17p deletion*', overlooking heterogeneity in prognosis owing to other clinical factors. Likewise, in GENETIC TESTING, he assumed that all reported *TP53* mutations confer an equally high risk of developing cancer.

A third way is in how he considered only limited types of evidence. In GENETIC TESTING, for example, Dr. Roy placed too much weight on the mechanistic link between *TP53* mutations and cancer, overlooking the

epidemiological evidence (or lack thereof) on the association between the particular variant and disease. Dr. Cong, by contrast, recognized the need to gather more evidence before reaching a conclusion.

A fourth way to describe Dr. Roy's errors is in terms of anchoring bias—the cognitive tendency to fixate on initial information and insufficiently adjust one's judgment in light of new data (Tversky and Kahneman 1975; Croskerry 2002). Anchoring may help explain the resident's failure to revise the prognosis he inferred from the presence of a *TP53* mutation in GENETIC TESTING or a 17p deletion in LEUKEMIA, even after learning that the former was a variant of uncertain significance and that the latter occurred in an asymptomatic patient with early-stage disease.[3]

All provide useful descriptions of how Dr. Roy erred. But what does each case have in common? Although the mistakes were made by a resident, they are not limited to novice doctors. Though presented as fictive dialogues, these examples illustrate a broader, empirically documented pattern in clinical reasoning: the tendency to give undue salience to certain types of information—particularly genomic biomarkers—in clinical decision-making (e.g., Lane et al. 2023; Callahan et al. 2022; Macklin et al. 2019). In each case, Dr. Roy placed disproportionate weight on a biomarker result. This led him to prioritize the biomarker defined reference class, overlook sources of heterogeneity within that reference class, and ignore other types of evidence. It is this undue salience, which explains Dr. Roy's faulty reasoning, that this article seeks to understand and correct. How did the consultant, Dr. Cong, avoid her resident's mistakes? In what follows, I argue that the difference lies in how each doctor treated biomarkers qua *abstractions* in clinical reasoning. Before delving into this difference, however, let's first take a closer look at biomarkers and their role in clinical practice.

2.2 What is a biomarker?

Biomarkers are biological data that serve as objective indicators of the state of a patient, where 'objective' is meant to contrast the more 'subjective' signs and symptoms gathered from clinical exam and history. Biomarkers have long been used in modern medicine, from measuring body temperature to detect febrile illness to serum creatinine to assess kidney function. Only more recently, however, have there been concerted efforts to clarify definitions to handle their growing numbers and applications.

Influential among these is the US Food and Drug Administration and National Institutes of Health Biomarker Working Group, who define a biomarker as 'an indicator of normal biological processes, pathogenic processes, or biological responses to an exposure or intervention, including therapeutic interventions' (2016, p. 45). They specify seven different categories of biomarkers based on their roles in research and medical product development.[4] For example,

---

[3] I thank an anonymous reviewer for suggesting this point, including the insight that anchoring based on biomarkers can influence not only diagnostic reasoning but also prognostic and therapeutic decision-making.

[4] These are diagnostic biomarkers, monitoring biomarkers, pharmacodynamic/response biomarkers, predictive biomarkers, prognostic biomarkers, safety biomarkers, and susceptibility/risk biomarkers.

predictive biomarkers identify the likelihood of favourable or unfavourable effects from exposure to a medical product (ibid., p. 19). Prognostic biomarkers, by contrast, identify the likelihood of a clinical event, such as disease progression, and are not defined in relation to a particular intervention (ibid., p. 23).[5] Predictive biomarkers play a key role in precision oncology, identifying individuals who may benefit from particular therapies, thus enabling 'personalized' treatment of cancer (Chin-Yee 2024). Each of the genomic biomarkers mentioned at the start—*5q deletion*, *NPM1 mutation*, *17p deletion*—fit one or more of the Working Group definitions.

While these definitions capture important uses of biomarkers, they remain geared towards uses in research and regulatory settings, i.e., for purposes of clinical trial design and pharmaceutical development, and stop short of examining how biomarkers are used in day-to-day clinical practice.

In clinical practice, biomarkers are tools to help resolve uncertain situations. This uncertainty arises in the context of patient care, the ultimate domain of biomarker application. In this context, practitioners often reserve the term 'biomarker' for a narrower class of indicators than those covered by the Working Group definitions, while at the same time including indicators that do not always fit neatly within a single category. In the clinic, calling something a biomarker confers a particular ontological status, indicating a link to a disease state or pathological process that holds special relevance within a given clinical context. In virtue of their special ontological status vis-à-vis disease, biomarkers often also have privileged epistemic status over other data, such as self-reported symptoms, in clinical reasoning.

Consider a common biomarker, hemoglobin A1c. Hemoglobin A1c serves as an indicator of dysfunctional glucose metabolism due to insulin resistance or deficiency. Its relation to what is taken to be the fundamental pathological process in diabetes confers upon this biomarker a privileged ontological and epistemic status in diabetes, where it is used to define the diagnosis, predict risk of specific complications, and establish treatment targets. Note that this privileged status might cause biomarker interpretation to go beyond the Working Group definitions, being used, for example, as a more general indicator of an individual's functional state. When a diabetes doctor approaches a patient with an elevated hemoglobin A1c, they seldom limit their interpretation to one of the seven Working Group definitions. Hemoglobin A1c is not only used as a prognostic biomarker, indicating higher likelihood of developing complications, but also (rightly or wrongly) to understand how the patient is managing and their overall state of health.

This example also highlights how a biomarker's salience is context dependent. Hemoglobin A1c is the biomarker of interest in the diabetes clinic. However, in follow up for the patient's prostate cancer, for instance, other

[5] Notice that in the Working Group definitions the role of biomarkers is stated in probabilistic, not causal, terms. Biomarkers indicate 'the likelihood of a clinical event' or 'individuals who are more likely' to experience favourable treatment response (2016, pp. 19-23) without necessarily playing a causal role in the clinical event or treatment response. Tabb and Lemoine (2021) pick up on this to distinguish between 'mechanistic' and 'statistical' biomarkers, with the former but not the latter indicating an established causal role. I return to this distinction in §3.1.

biomarkers, such as prostate specific antigen (a tumour marker), become more salient to clinical reasoning. In this new context, hemoglobin A1c might be ignored or treated as just another variable, alongside cholesterol or blood pressure, not serving as relevant biomarkers in the situation at hand. Biomarkers must therefore be defined in relation to a particular clinical context or question.[6] As LEUKEMIA illustrates (discussed further below), this clinical context includes not only the type of disease, but also the specific stage and treatment decision a patient faces.

Technologies such as high-throughput analyzers have expanded use of biomarkers in practice, generating rapid data outputs for use in clinical decision-making. Such developments spurred the growth of genomic biomarkers in particular, which measure genes or effects of genes using genomic sequencing technologies. Genomic biomarkers increasingly impact decision-making in cancer medicine, from influencing choice of therapy in solid tumours, such as melanoma, breast, and lung cancer, to defining subtypes and guiding treatment of blood cancers, such as acute and chronic leukemias. In cancer especially, genomic biomarkers are often afforded special significance due to the dominant view of cancer as a genetic disease.[7] In clinical practice, such views make genomic biomarkers particularly amenable to being interpreted in mechanistic or causal terms, despite the Working Group definitions remaining agnostic about whether or not biomarkers *cause*, or are merely associated with, particular outcomes (see Tabb and Lemoine 2021). We saw this, for example, in how the resident in LEUKEMIA interpreted the biomarker, the 17p deletion, as *causing* rapid disease progression. I return to this point below. Given precision oncology's emphasis on genomics, the remainder of this article focuses on genomic biomarkers, although the ideas apply to biomarkers more generally.

2.3 Genomic biomarkers and epistemic harms

While genomic biomarkers promise to usher in an era of personalized medicine they are not without their epistemic challenges. The cases from §2.1 illustrate how errors arise when genomic biomarkers are given undue salience in clinical reasoning. If unchecked, such errors can easily harm patients. One dimension of this harm is clearly prudential—for example, if Dr. Roy the resident exposes Mrs. de Vries, his patient in THROMBOSIS, to unnecessary treatment, or causes Ms. Ali, his patient in GENETIC TESTING, distress from a mistaken diagnosis and prognosis.

But another dimension of this harm is epistemic. Inapt categorization and risk communication can also harm patients in their capacities as knowers. Or so some philosophers have claimed when it comes to genomic biomarkers. Reynolds (2020), for example, argues that genomic technologies contribute to epistemic harms through the way that their products, i.e., genomic biomarkers, are

[6] 'Context of use' is discussed in the Working Group definitions but primarily refers to research and regulatory contexts (FDA-NIH Biomarker Working Group 2016, 48). See also Leptak et al. (2017) for discussion of (regulatory) context of use of biomarkers.

[7] See, for example, www.cancer.gov/about-cancer/causes-prevention/genetics. See also, for discussion, Tabery (2024).

interpreted in the clinical encounter. They highlight two mechanisms behind these epistemic harms:

(i) *Epistemic capture:* 'when fundamentally ambiguous information is transformed into and treated as definite forms of knowledge'; and
(ii) *Value partitioning:* 'when variously available hermeneutical resources are restricted to an inappropriately simplistic evaluative scale or set of appraisals' (p. 171).

Reynolds contends that (i) and (ii) are distinct ways that epistemic harms result from genomic technologies, which present uncertain information in certain terms and simplistically delineate the normal from the pathological.

Reynolds's motivating cases are similar to GENETIC TESTING, focusing on how *variants of uncertain significance* are communicated to patients. The worry is that, like our resident Dr. Roy in GENETIC TESTING, clinicians tend to treat ambiguous genomic information as definitive ('*causes leukemia*') and evaluate it in simplistic terms ('*this abnormality*') in a way that excludes other sources of evidence. As a result, patients are prevented from bringing other interpretive resources—such as their own lived experience—to bear on the clinical encounter. Reynolds sees this as an important way that genomic technologies harm patients in their capacities as knowers.

Some might find Reynolds's worries overstated. Firstly, one might question whether there is, in fact, epistemic harm in such cases beyond the prudential harms that result from wrongly communicating genomic information, i.e., wrongly treating uncertain variants as certain. In GENETIC TESTING, for example, Ms. Ali might be misled into believing she will develop leukemia, causing her undue anxiety. But is she really harmed *epistemically*, i.e., in her capacity as a knower? Reynolds would argue yes: Ms. Ali is epistemically harmed insofar as genetic testing is used to exclude or dismiss other sources of knowledge about her health, such as her own bodily experiences or personal understanding of her well-being. If the genomic biomarker, despite its uncertainty, prevents Ms. Ali from recognising herself as healthy, she suffers an epistemic harm: her ability to understand her own state of health is undermined, and her lived experience as a credible source of knowledge is dismissed.

However, even if one accepts Reynolds's argument that such cases involve epistemic harm, say, by marginalising patients' lived experiences, it's clear that these harms do not necessarily follow from the use of genomic technologies. Rather, they are contingent on several factors, not least on how genomic information is communicated in the clinical encounter. In GENETIC TESTING, for example, epistemic harms may be reduced or even eliminated if, say, Dr. Cong, the cautious consultant, rather than Dr. Roy, the rash resident, communicated the results. The situation may be further helped by a dedicated genetic counsellor (Makhnoon et al. 2021). Indeed, research on genetic counselling for hereditary cancers shows that how variants of uncertain significance are communicated has a significant impact on patients' perceptions, including whether they view results as pathological (Makhnoon et al. 2018; Zhong et al. 2021). Many patients, in fact, can appreciate ambiguity and resist interpreting genomics as conclusive (Ishak et al. 2023). Moreover, at least in some cases, genomic technologies can complement, rather than replace, other resources that patients use to understand

their risk, such as knowledge of their family history or lifestyle (Scherr et al. 2021). Genomic technologies, it seems, do not inevitably cause epistemic harms as Reynolds seems to suggest.

Moreover, one might doubt whether Reynolds's concerns extend to genomic biomarkers more generally, beyond variants of uncertain significance which are—by definition—ambiguous in their clinical significance.[8] Variants of uncertain significance indeed pose epistemic problems for precision medicine (Tonelli and Shirts 2018). But not all genomic information is 'fundamentally ambiguous' (Reynolds 2020, p. 162). Nor is all 'ambiguous' genomic information ambiguous in the same way.[9] In an effort to reduce ambiguity, guidelines for reporting genetic variants have been proposed and widely adopted (Li et al. 2017). Such guidelines include a type of evidence hierarchy specifying tiers of increasing clinical significance for genetic variants, moving from 'likely benign' to 'unknown' (i.e., *variants of uncertain significance*) to 'potential' to 'strong', based on several factors, including study size, reproducibility, expert consensus, and existence of approved treatments for a given variant. Likewise, a multitude of systems for ranking clinical 'actionability' of genomic information have been devised (see Chin-Yee and Plutynski 2023). While the merits and demerits of these different approaches can be debated, at the very least they show how some genomic information is treated as less ambiguous in its clinical significance. Claims on the basis of variants of 'strong' clinical significance or 'highly actionable' variants are, plausibly, less ambiguous; Reynolds's worry about epistemic capture, which is premised on genomic information being 'fundamentally ambiguous', therefore, may not apply.

Though one might question the overall framing, necessity, and scope of Reynolds's concerns, they nevertheless provide a useful starting point for thinking about the harms posed by use of genomic biomarkers in clinical practice. While Reynolds's worries may not generalize to all genomic biomarkers, across all clinical encounters, they begin to expose an issue that should be taken seriously in precision medicine and oncology—namely, that there is always a risk of harm akin to epistemic capture in applying genomic biomarkers in practice. Left uncorrected, the resident in THROMBOSIS, GENETIC TESTING and LEUKEMIA does, in fact, run the risk of inflicting harm to his patients, and this harm can have an epistemic dimension.

How can we better come to grips with this harm? Not only, *pace* Reynolds, by pointing upstream to the types of 'epistemic frameworks' at play with genomic technologies (2020, p. 163). Nor simply by looking downstream at (mis)communication of genomic information in token clinical encounters. Each perspective alone is insufficient to capture the problem, which takes root somewhere in between. It takes root, I argue, in how biomarkers are used as *abstractions* in clinical reasoning.

---

[8] Though Reynolds focuses on variants of uncertain significance, they take these epistemic harms to be a more general feature of genomic technologies (2020, p. 162).

[9] For example, beyond ambiguities in sequence interpretation and 'variant calling', uncertainty still remains with known pathogenic variants, which at best give probabilistic information on the likelihood of a disease phenotype owing to incomplete penetrance of most genetic variants associated with human diseases.

2.4 Abstraction and reification

I said in §2.2 that, from the perspective of clinical practice, biomarkers are tools to help resolve uncertain situations. In the utterance—*This is a case of chronic lymphocytic leukemia with 17p deletion*—the biomarker conveys information intended to help determine prognosis and treatment. The utterance highlights what is taken to be most salient for clinical reasoning, while omitting all other features from the actual case. In this way, it is an *abstraction*.[10]

Abstraction is a central topic in philosophy of science, and to understand the role of abstraction in clinical reasoning it helps to consider the analogy with scientific reasoning. Here I draw on Winther's (2014, 2020) recent discussion of abstraction in the sciences, which offers important insights for medicine that I begin to unpack here.

Winther (2020) focuses on maps and mapmaking—not only literal maps, i.e., cartographic representations, but also metaphorical maps, which are common across the sciences. Consider, for example, The Cancer Genome *Atlas* Project, a large-scale initiative from the US National Cancer Institute. This project sequenced samples from over 11,000 patients with thirty-three different tumour types over a twelve-year period in an endeavour to 'map' the cancer genome and 'chart a new course across the complex landscape of human malignancies' (Collins and Barker 2007). Cancer medicine is suffused with map metaphors, with advances in sequencing technologies offering new methods for *mapping* genomic *landscapes* to *guide* research and practice.

Winther details how mapmaking, in moving from the world to a representation, involves abstraction. In cartography, for example, abstraction occurs in multiple steps, from processes calibration and measurement to stages of scaling, simplification and symbolizing. Each step incorporates assumptions which shape the abstract representation. Analogous processes were at play in the mapmaking endeavours of The Cancer Genome Atlas Project, a massive enterprise in abstraction, starting from thousands of human subjects with cancer

[10] One might observe that biomarker abstractions resemble what Wasserman and Asch (2005) call *synecdoche*—a process where a particular stigmatized trait comes to stand in for, and thereby 'obscures or effaces', the whole person (p. 173). On their account, the 'sin of the synecdoche' refers to a moral failure arising when a single, often devalued, trait comes to dominate how a person is regarded as a whole. More generally, however, in literary analysis and linguistics, *synecdoche* refers to a morally neutral rhetorical device—a type of metonymy in which a part stands in for the whole. In this broader sense, many clinical representations can take a synecdochic form: a biomarker or diagnosis (e.g., 'the pneumonia') may come to stand in for a patient's entire situation. However, *clinical abstraction*, as I use the term, differs from both senses of synecdoche. Whereas synecdoche functions referentially, clinical abstraction involves a predicative move: it attributes a property to a patient rather than substituting that property for the person as a whole. Moreover, while clinical abstraction, like synecdoche, involves omission, it need not entail the kind of epistemic distortion or moral failing associated with the 'sin of the synecdoche' (ibid., p. 181). As I argue in what follows, abstraction is a routine and often epistemically productive part of clinical reasoning, where foregrounding certain clinical features (including biomarkers) helps guide decision-making. I thank an anonymous review for suggesting the relevance of synecdoche to this discussion.

to produce representations of their disease biology in genomic terms.[11] While such projects represent important advances in cancer research, these mapmaking endeavours must be paired with an awareness of the challenges posed by abstraction, especially when translating findings into practice.

This is a major concern for Winther, who worries about the ways abstraction can go awry. His main worry is *reification*, which occurs when an abstract representation is treated as a concrete entity, taken to fully describe and even constitute the target phenomenon in its entirety (Winther 2020, p. 89).[12] This overestimates a representation's scope and efficacy, leading it to be applied beyond its appropriate context. Reification thus arises from how abstractions are made to do work in the world as a basis for inference and action.

Sticking to cartography, consider the Mercator projection, a representation of the globe initially produced in 1569 for the purposes of navigation but which has since become the most common depiction of 'the world map'. The Mercator projection inflates size away from the equator, leading it to exaggerate, for example, the land mass of Greenland relative to Africa. (It shows them as roughly the same size, whereas Africa is in fact is fourteen times larger than Greenland in surface area.) The Mercator projection served as a useful tool for early modern navigators. But its persistent use as a world map is an instance of reification. Here, reification arises when the Mercator projection qua abstraction is put to use in a way that forgets the projection's history and intended function, applying it beyond its appropriate domain. The student who, treating the Mercator projection as the 'the world map', infers that Africa is the same size of Greenland, has *reified* this representation. Reification—taking this map to capture all or most geographic facts about the world—underlies the student's faulty inference. Science is rife with reifications, from constructs of race and intelligence to selfish genes.[13]

What leads us to cross the line from abstraction to reification? Winther highlights two factors that contribute to reification, *universalizing* and *narrowing*:

> 1. *Universalizing* is holding your map, theory, model, or other representation or abstraction to encompass all, or almost all, phenomena within a given domain… What falls outside of this scope is considered irrelevant or uninteresting, or both…
> 2. *Narrowing* is constraining theoretical content to a few central laws, principles, and assumptions, thereby diminishing the nuance and diversity of models and theoretical components. The internal

[11] Plutynski (forthcoming) offers a detailed study of the range of assumptions that informed TCGA, including considerations in data analysis, such as how to classify tumours and define 'driver mutations'. Such assumptions in mapmaking all shaped the view of human cancer afforded by its projects. The cancer genome 'maps' produced had a profound impact on current understandings of human malignancies, largely bolstering the dominant model of cancer as a genetic disease and promoting approaches to classification, prognostication, and treatment in genomic terms.

[12] Problems of reification have been raised by several philosophers, most notably in the pragmatist tradition, variously described as 'vicious abstractionism' or 'the psychologist's fallacy' (James), 'misplaced concreteness' (Whitehead), and 'hypostatization' or 'the philosophical fallacy' (Dewey). See, for discussion, Winther (2014).

[13] For discussion on reification of race, see Kaplan and Winther (2013); on intelligence, see Gould (2006); on genes, see Walsh (2020).

> heterogeneity of theory decreases, while explanatory power might seemingly grow… (2020, p. 91).

To illustrate these factors, consider an example from cancer medicine involving the use of genomic biomarkers for prognostication. Myeloproliferative neoplasms are a group of blood cancers that have seen major clinical advances over the past decade in large part thanks to landmark genomic studies (see Spivak 2017). Prognostic scoring systems have long been used to predict survival in these diseases and now increasingly incorporate genomic biomarkers. For example, in myelofibrosis, a type of myeloproliferative neoplasm, the International Prognostic Scoring System was previously the 'gold standard' for prognostication (Cervantes et al. 2009). It predicts survival based on five variables, including whether the patient has constitutional symptoms, which refer to weight loss, fevers, and drenching night sweats. More recently, however, it has been supplanted by newer scoring systems. The Genetically Inspired Prognostic Scoring System (GIPSS), for example, seeks to 'limit prognostic reliance on clinical variables' by predicting survival based solely on genetic mutations (Tefferi et al. 2018).

As tools, these prognostic scores can prove useful, abstracting away from clinical complexity to highlight, for instance, how particular genetic mutations might influence a patient's survival. But danger arises from universalizing these representations—from mistaking this limited 'map' of genomic biomarkers as constituting *all* the relevant data, rendering all other factors within the clinical context irrelevant or uninteresting. There is also a risk of narrowing here, which occurs if application of the score causes the clinician to lose sight of nuance and heterogeneity within each genetic subgroup. Together, universalizing and narrowing could lead the cutting-edge oncologist to *only consider* genetic mutations in the newer scoring system in their clinical reasoning. Thus, when seeing a patient with myelofibrosis, say, Ms. Ma, they risk interpreting her symptoms of weight loss and drenching night sweats as *irrelevant* to her prognosis—or worse, to her warranting treatment. They mistake the map for the territory.

To be clear, if the cutting-edge oncologist is purely after predicting survival, they might be well-served by this newer scoring system. It might very well be the case that leaving out symptoms and considering only genes improves survival prediction in the population relevant to their practice. This is, of course, an empirical question. But as soon as the oncologist takes the score to offer a more global assessment of their patient—to offer a guide to *treating Ms. Ma*—reification rears its head.[14] Like the resident Dr. Roy in LEUKEMIA, the cutting-edge oncologist must be careful describing Ms. Ma as *'the case of myelofibrosis with high-risk GIPSS*' and basing further inferences on this description alone. While this abstraction may say something quite specific about Ms. Ma's genomic biomarkers and, applied in the right context, her expected survival, it leaves out many other factors essential to her care. These factors range from her clinical symptoms of weight loss and night sweats that may require supportive therapies, to her social

[14] Moreover, such scores employ a narrow definition of prognostication as mere prediction, which jettisons a richer spectrum of activity involved in prognostication in clinical practice (see Thomas et al. 2019). This is a second way use of such scores could exhibit 'narrowing' in Winther's sense.

circumstances that could determine whether certain treatments are feasible or appropriate in the first place.

Of course, the formal purpose of a tool like the GIPSS is to predict survival, *not* to guide comprehensive treatment decisions.[15] But when the oncologist forgets this purpose and treats this genomic score as a comprehensive representation of the patient—that is, when they *reify* the abstraction—they risk overlooking key aspects of the case and providing suboptimal care. This use of the score forgets that, in practice, clinical variables like a patient's symptoms are not mere 'surrogates' for 'underlying genetic lesions' (Tefferi et al. 2018, p. 1640); rather, they are essential guides to patient care. Symptoms like Ms. Ma's might signal disease progression, prompt clinical intervention, or warrant treatment on their own. The GIPSS score, qua abstraction, offers only a narrow glimpse into Ms. Ma's condition.

This example illustrates how universalising and narrowing, if unchecked, can lead to reification of a given abstraction. Reification, in turn, gives that abstraction undue salience in clinical reasoning, causing practitioners to ignore other relevant features of the situation from which it is abstracted. It is reification that underlies Dr. Roy's faulty reasoning with genomic biomarkers in THROMBOSIS, GENETIC TESTING, and LEUKEMIA.

Reification also underlies the epistemic harms that concern Reynolds (2020) surrounding genomic technologies. Where Reynolds sees epistemic capture and value partitioning, we can now see that genomic information is *universalized*, put forward as the privileged interpretive resource to the exclusion of other ways of knowing, and *narrowed*, presented along overly simplistic scales in a manner that conceals ambiguity and nuance. Moreover, in Reynolds's examples, genomic technologies are applied without proper attention to their scope and intended function, i.e., the purpose of clinical genetic testing is to identify specific, 'actionable' mutations, not variants of uncertain significance.

This reframing, which sees biomarker use through the lens of abstraction, allows us to recognize these epistemic harms as results of reification. Unlike Reynolds's account, however, this view also allows us to appreciate the epistemic benefits of biomarkers qua abstractions in clinical reasoning. We are now better positioned to again ask the question from §2.1—how did the consultant, Dr. Cong, avoid her resident's mistakes? She did so by exemplifying an approach that successfully guards against reification but nevertheless preserves the centrality of abstraction in clinical reasoning. I call this approach *clinical abstraction*.

3 Clinical abstraction

3.1 Three maxims of clinical abstraction

[15] Though one might easily be misled into believing it's the latter given the way such scores or biomarkers are presented in treatment algorithms as key determinants of clinical decisions, effectively suggesting they *are* comprehensive guides to treatment. See, *inter alia*, Tefferi et al. (2018, Figure 5); see also Nabhan et al. (2015, Figure 2) for an example from chronic lymphocytic leukemia.

What is clinical abstraction? Clinical abstraction is a normative account of how doctors ought to reason with abstractions. It's also, I contend, a descriptive account of what many experienced doctors already do. This section outlines the key features of clinical abstraction—that is, *good* clinical abstraction—and how it guards against bad uses of abstractions in medicine.

I should emphasize at the outset that clinical abstraction is *not* proposed as a rival account of scientific abstraction. Rather, it is an approach to abstraction oriented towards the practical aims of clinical medicine, where decisions must be made under uncertainty in service of individual patients. While doctors and scientists may use abstraction in similar ways—for example, to simplify complexity or highlight salient causal features—clinical abstraction is guided by the need to make judgments and take action in the care of particular persons.

This need not imply a sharp boundary between science and medicine. Applied scientific domains—such as meteorology, seismology, or conservation biology—also rely on abstractions to support risk management under uncertainty, including decisions about how to act in specific cases.[16] When skillfully applied, abstraction in these fields may share features with clinical abstraction—an approach attuned to the ethical and epistemic demands of bedside medicine.

At the bedside, abstraction is ultimately tied to action. It helps clinicians think and act, here and now, for a particular patient. This contrasts with several leading accounts of scientific abstraction, where abstraction involves selective omission to isolate features of causal or explanatory relevance.[17] While clinical abstraction shares the idea that abstraction involves omission of certain features to highlight others of special relevance, for clinical abstraction, relevance is not tied to some broader explanatory aim but rather to action in the individual case.[18] This action might take the form of deciding a treatment, communicating a prognosis, or ordering more tests—each directed at the care of a particular individual. From this, we can distill what I call the first maxim of clinical abstraction:

> *Maxim one.* Perform abstraction with the express purpose of guiding thought and action in the face of a particular uncertain situation.

According to this maxim, one should engage in clinical abstraction—for example, representing a patient in terms of their genomic biomarker—only insofar as it serves one's action in the particular situation at hand. If it doesn't, the clinician should set aside that representation for the time being. In THROMBOSIS, for example, moving from Mrs. de Vries, who suffered a blood clot after hip surgery, to '*a case of deep vein thrombosis in a Factor V Leiden heterozygote*' does not help reduce uncertainty and guide action in her case, at least when it comes to deciding initial treatment. For this reason, the consultant, Dr. Cong, set aside this abstraction in favour of another one, based on Mrs. de Vries's post-surgical status, to help determine treatment.

---

[16] I thank an anonymous reviewer for suggesting this point.

[17] See, *inter alia*, Cartwright (1984, 1994), Jones (2005), Godfrey-Smith (2009), Strevens (2008), Potochnik (2017), and Levy (2018).

[18] This is not to say that doctors are not sometimes after explanation in their uses of abstraction; however, clinical abstraction seeks explanation insofar as it enables action in the individual case.

Both abstractions, however, would be licensed on accounts of scientific abstraction that emphasize abstraction's role in explanation or understanding. A scientist can retain both representations so long as they offer fruitful paths to understanding the broader phenomenon in question, i.e., blood clot recurrence *in general*. But for Dr. Cong, the need care for a *particular* patient with a clot requires selecting the abstraction that best guides her treatment.

This is not to say that clinical abstraction stands apart from scientific abstraction. Dr. Cong's judgment relied on researchers having pursued both abstractions, i.e., by conducting the relevant cohort studies, to know that one abstraction better predicts the risk of recurrent blood clots in her patient Mrs. de Vries. To reiterate, clinical abstraction is best conceived not as an alternative to scientific abstraction but rather as an extension to a domain with particular aims and constraints.

To return to the map analogy, doctors are more akin to navigators than cartographers. A navigator may follow their oceanographic map to capitalize on a favourable current, or their weather map to avoid an impending storm. But when the two suggest divergent courses, the navigator must choose. They cannot steer port and starboard simultaneously. Facing such choices, they must not rely too heavily on any single map. As any experienced navigator knows, a map is not the territory. [19]

Ditto for doctors. By conceiving of abstraction as a tool for thought and action in a particular situation, clinical abstraction maintains contact with the territory—with the concrete situation from which the abstraction was derived. This helps prevent the doctor from conferring upon the abstraction an independent standing that belongs only to the concrete situation. For the cutting-edge oncologist, for example, seeing Ms. Ma in clinic is a potent reminder that '*the case of myelofibrosis with high-risk GIPSS*' only exists by abstracting away from Ms. Ma, the patient.

Clinical abstraction's grounding in the concrete case serves as an initial defence against reification. However, given the known harms of reification in practice, clinical abstraction goes one step further to guard against them. How? By deliberately *questioning the omissions* required to produce an abstraction. Think back to LEUKEMIA, for example. There, the consultant Dr. Cong recognized that her resident's abstraction involved omission and selective emphasis. Asking the question—*what is left out?*—allowed Dr. Cong to correct her resident. This question highlighted crucial clinical features that Dr. Roy omitted, which in turn led him to make untenable inferences about the patient's prognosis and treatment. In clinical abstraction, this question—*what is left out*?—is not raised merely to interrogate representational truth, or simply to determine the extent to which a given abstraction (mis)describes the actual case. Rather, it is raised to scrutinize an abstraction's usefulness in the situation at hand; to ask—*how do the omissions impact application?* This is the central question faced in clinical practice. From this, we can distill a further two maxims of clinical abstraction:

[19] 'A map is not the territory' said Korzybski, quoted in Winther (2020, p. 43).

> *Maxim two.* Avoid reification, that is, do not treat abstractions as entities with independent standing beyond that of tools for thought and action in a concrete situation.
>
> *Maxim three.* Question the omissions that produce an abstraction and seek to understand how these omissions impact application in the situation at hand.

These two maxims, to be sure, express considerations not entirely unique to clinical abstraction. Scientists, too, should guard against reification. This is especially true in applied fields, such as engineering, epidemiology, and environmental science, where overlooking an abstraction's omissions can lead to misguided decisions and real-world harms. Still, the risks of reification are particularly acute in clinical medicine, where abstractions are routinely brought to bear on concrete cases, and where, as we have seen, they pose distinct prudential and epistemic harms if reified. Maxims two and three, therefore, reflect added measures taken by clinical abstraction to prevent these harms.

To briefly recap, then, I've outlined three main features of clinical abstraction. First, it sees abstraction as a tool to guide thought and action in a particular case, rather to serve broader aims such as explanation. Second, it emphasizes the need to avoid reification because of the harms it poses to patients. Third, it does this by deliberately scrutinizing the omissions behind a given abstraction and how they impact that abstraction's usefulness in a particular case.

At this point, some might worry that clinical abstraction underemphasizes abstraction's explanatory role, leaving out many important uses of abstraction in medicine. Specifically, one might worry that explanation serves as an indispensable intermediary between abstraction and action in clinical medicine. For example, although the function of a biomarker qua abstraction may be to indicate a specific action (e.g., prognostic, therapeutic, etc.), this action is only justified on the basis of an explanation (e.g., the biomarker serves *as an explanation for* poor prognosis or favourable response to a drug). Moreover, this explanation might refer to causal capacities isolated by a given abstraction, for instance, highlighting a certain genetic mutation's causal role in disease progression or a drug's mechanism of action.

To be clear, clinical abstraction does not dismiss the importance of causal explanation in medicine. Causal explanation is often an important desideratum for abstractions used to justify clinical action. Yet it is not the sole, nor even the ultimate, aim of abstraction in this domain. For example, abstracting from a sick patient to '*a case of fever with blood cultures growing gram-positive cocci in clusters*', may provide a causal explanation—namely *S. aureus* bacteremia causing sepsis—of the patient's illness. But the abstraction fails to provide a causal explanation if, say, the patient's fever is instead caused by influenza and the positive blood cultures by an incidental skin contaminant (a common scenario). Facing a sick patient, however, in both cases the doctor may use this abstraction to guide action—i.e., to give empiric antibiotics—regardless of whether the abstraction isolates the cause of the patient's illness. Seeing abstraction as a mere means to causal explanation fails to capture why this abstraction proves useful to the doctor—even to the doctor who suspects her patient has influenza. It fails to fully capture abstraction's uses in clinical practice to guide action in uncertain situations—

situations where causal knowledge is often unavailable, nor always necessary to warrant interventions with proven empirical success. Clinical abstraction can accommodate such uses, whereas accounts that tolerate abstraction only insofar as it supports causal explanation cannot.

This is important to bear in mind with abstractions involving biomarkers, as interpreting them as giving causal explanations can be misleading. Though Tabb and Lemoine (2021, p. 194) attempt to distinguish between 'mechanistic' and 'statistical' biomarkers, with only the former having a causal role, this distinction is not commonly made in clinical practice. Despite avoidance of causal language in the Working Group definitions, in practice there's a tendency to attribute causal powers to biomarkers, in turn contributing to their epistemic privileging. This is especially the case with genomic biomarkers in cancer, where genetic mutations are not only assigned causal roles but are often regarded as the primary drivers of disease progression and, therefore, as the most promising therapeutic targets. We saw this at play in LEUKEMIA, for example, where the resident, Dr. Roy, takes the genomic biomarker, 17p deletion, as directly causing rapid disease progression and therefore warranting targeted treatment. Among practitioners, Dr. Roy is not alone in this interpretation. Precision medicine and oncology tend to emphasize genetic causation (Tabery 2024), following a longer history of genetic exceptionalism in medicine (Garrison et al. 2019, Boyce 2019). In practice, these factors make genomic biomarkers especially prone to reification and its associated harms.

In short, while clinical abstraction does not deny that abstractions can serve causal or explanatory functions in medicine, it recognizes these functions as part of fallible toolkit whose primary aim is to guide action in the face of a particular uncertain situation.

Here, one might also question whether clinical abstraction, as I've just described it, is really necessary. What do its maxims add beyond other guides to medical decision-making? For example, it may be true that Dr. Roy violated clinical abstraction's three maxims in how he relied on abstractions that were not appropriately action guiding, treated them as concrete, and failed to consider their omissions. But one could easily point to other ways that his decision-making erred without reference to these maxims. I already alluded to some of these in §2.1. For instance, in labelling Ms. de Vries in THROMBOSIS as '*a case of deep vein thrombosis in a Factor V Leiden heterozygote*' to determine recurrence risk and recommend treatment, Dr. Roy's choice of reference class was not the 'best reference class available', such as the 'narrowest reference class' for which there is reliable statistics (Reichenbach 1949, p. 18), or the 'broadest homogenous reference class' to which the case belongs (Salmon 1971, p. 43), to cite just two classic ways to address the reference class problem.[20] A better reference class was instead supplied by the consultant Dr. Cong based on Ms. de Vries's post-surgical status and reliable evidence of recurrence risk in this population.[21] More generally, one

[20] These, of course, do not *solve* the reference class problem (see Hájek 2007), though these approaches may have helped prevent Dr. Roy's error in THROMBOSIS.

[21] Note, there is lack of reliable statistics for recurrence risk in patients with thrombosis after major surgery *and* who are heterozygous for *Factor V Leiden*; however, if available, presumably choosing a reference class defined by both the biomarker and post-surgical status would be best.

could simply point out how Dr. Roy failed to consider the 'total evidence' in the case, for example, ignoring crucial clinical information that should have factored into his diagnosis and management.

These are, to be sure, useful ways to describe such errors in reasoning, and suggest correctives without recourse to clinical abstraction and its maxims. A single, alternate principle—say, *always consider the total evidence*—might seem sufficient to set Dr. Roy straight. Clinical abstraction, however, has the advantage of providing Dr. Roy with more specific guidance in such cases. It does so by answering more fine-grained, prior questions that are not immediately addressed by the other descriptions. Namely, *why* did Dr. Roy gravitate towards the biomarker defined reference class? *Why* did he systematically ignore other evidence in each of the three cases? My analysis answers these *why* questions by pointing to problems of abstraction and reification. While strict adherence to a more general principle of total evidence might also avoid Dr. Roy's errors, it lacks the specificity to guide his actions in the cases. There are a great many ways that one can fail to consider the total evidence. Therefore, instead of resting on this general principle, it helps to look at more specific ways it is instantiated, or fails to be instantiated, in clinical reasoning.[22] It helps to look at the processes that gave biomarkers undue salience in Dr. Roy's reasoning, which led him to ignore other evidence in the first place.

Clinical abstraction is, of course, not the only lens through which to view errors in clinical reasoning. Nor does it necessarily provide a comprehensive picture of such errors. However, it does help to capture—with a clear, corrective focus—what I take to be a common and remediable source of faulty reasoning with biomarkers in clinical practice.

3.2 Clinical abstraction in practice

Clinical abstraction not only offers a normative guide for how doctors should reason with abstractions but also, as the cases in §2.1 illustrate, a descriptive account of what many experienced doctors already do in practice.

Clinical abstraction, in fact, closely parallels how clinical reasoning is taught in health professions education. For example, so-called *illness scripts* are a common device used in medical training, which take the form of a narrative designed to highlight salient features from a clinical scenario in order to prompt specific diagnostic and treatment considerations.[23] For example, the 'young adult with weight loss, drenching night sweats and enlarged lymph nodes, painful with alcohol consumption' alerts the trainee to a diagnosis of Hodgkin lymphoma, a cancer of the lymph glands, with a differential diagnosis that includes infection or other cancer. The 'elderly patient with deep vein thrombosis following hip replacement' indicates to the trainee that three months of blood thinners is

[22] Though some might interpret clinical abstraction and its maxims along these lines, as a specific instantiation of a more general principle of total evidence, others might instead wish to interpret these maxims as practical heuristics for medical decision-making under non-ideal conditions, where constraints on time, cognitive resources, evidence availability pose further challenges to applying a principle of total evidence. I leave this open as another possible interpretation of my account.

[23] For discussion of illness scripts in medical education, see Custers (2015).

generally sufficient. These illness scripts are abstractions, meant to represent the 'textbook' case presentation for a given disease. Cases often appear on licensing exams in the same format as these illness scripts, testing candidates' ability to reason with abstractions.

The bulk of medical training, however, is experiential. This is where trainees, like our resident Dr. Roy, must learn through trial and error to abstract from concrete cases, select the right abstractions from a set of shared representations, and appropriately apply them in actual cases. Such is a core competency of medical training. This competency requires learning not only how to abstract but also how to de-idealize. It requires learning, like the experienced navigator, the dangers of excessive reliance on one set of abstractions. Dismayed by his error in LEUKEMIA, Dr. Roy learns to temper his reliance on the biomarker abstraction and attend to other features of the case.

The centrality of abstraction in medicine extends beyond education. It is seen throughout many aspects of routine practice. For example, in clinical rounds, where health professionals meet to discuss patients, case presentations have the same structure as illness scripts. In cancer medicine, these rounds are referred to as tumour boards. Precision oncology has spawned a new type of tumour boards—so-called *molecular* tumour boards—focused on genomic biomarkers (Chin-Yee and Plutynski 2023). Presentations at tumour boards, molecular or otherwise, typically lead with sentences like the ones we saw at the outset. *This is a case of myelodysplastic syndrome with 5q deletion. This is a case of acute myeloid leukemia with NMP1 mutation.* Here, abstraction facilitates communication between professionals and allows collective reasoning about a case. In this context, abstraction serves as a powerful tool to enable more complex, higher order reasoning, removed from the messiness of the clinical encounter. However, as the experienced oncologist recognizes, all of this abstract discussion serves one goal: to help guide action in the particular case. In light of this goal, recommendations made at tumour boards are just that—recommendations. These must always be tailored—that is to say, de-idealized—to the specific circumstances of each patient.

To better illustrate how clinical abstraction addresses the epistemic challenges raised in §2, let's now return to one of the cases. *This is a case of chronic lymphocytic leukemia with 17p deletion.* This case, the basis for the fictive dialogue in LEUKEMIA, was in fact inspired by a real patient scenario.[24] However, the actual scenario did not involve the resident, Dr. Roy, nor (unfortunately) the consultant, Dr. Cong, who corrected him. Instead, it involved a different doctor. Let's call him Dr. Hunt.

The patient, Mr. Lund, a seventy-year-old retired teacher, was referred to Dr. Hunt, a hematologist, for an elevated white blood cell count noted on routine bloodwork. Mr. Lund was otherwise healthy and asymptomatic. Dr. Hunt ordered the standard diagnostic tests, which confirmed the diagnosis of chronic lymphocytic leukemia (CLL). CLL is, generally speaking, a slow-growing type of leukemia and one of the most common blood cancers in the Western world. Given that Mr. Lund was feeling well, with no other symptoms or abnormalities

[24] This scenario comes from the author's clinical experience with all details altered or removed to protect confidentiality.

on his physical exam and blood tests, he was classified as having early-stage disease, for which the recommended management is active surveillance without any treatment. This is sometimes called a 'watch and wait' approach. This recommendation is based on early-stage CLL having a good prognosis, a median survival of over ten years, with one third of patients Mr. Lund's age never requiring treatment in their lifetimes.

However, in this particular case, both Mr. Lund and his hematologist Dr. Hunt desired a more precise prediction, beyond these general numbers for allcomers with early-stage disease as measured in large cohort studies. Dr. Hunt wondered—what is *this individual's* risk of disease progression? As we've seen, this is a key question in precision oncology, which aims to personalize prediction and treatment decisions based on biomarkers. To this end, Dr. Hunt applied the (then state-of-the-art) genomic biomarkers, in this case cytogenetic testing for chromosomal abnormalities found in CLL. This testing revealed the 17p deletion.

17p, the short arm of chromosome 17, is the locus of the tumour suppressor gene *TP53*, which plays a key role in DNA repair and cell death, negatively regulating cell proliferation and survival. In CLL, loss of 17p (i.e., 17p deletion) serves both as a prognostic and a predictive biomarker. It is associated with more rapid disease progression and reduced survival. A landmark study from the early 2000s found that patients with 17p deletion had survival times around one to two years versus ten to twelve years for patients with normal cytogenetics (Döhner et al. 2000). More recently, 17p deletion has become a predictive biomarker in CLL, associated with favourable response to newer drugs compared to conventional chemotherapy (Woyach et al. 2018). There is also mechanistic rationale behind this biomarker, which is interpreted as playing a causal role in disease progression through loss of expression of the tumour suppressor gene *TP53*.

Dr. Hunt now classified Mr. Lund as *a case of CLL with 17p deletion*. He told Mr. Lund that his CLL harboured an adverse genomic biomarker. He told him that this meant a far worse prognosis than initially thought, reducing his life expectancy by around a decade. Mr. Lund, who continued to feel well and comparatively healthy, was now told that he had 'bad disease'. He continued on a watch and wait approach. But now this came with significant anxiety—was he facing the last years of life? The biomarker seemed to suggest as much. Dr. Hunt, searching for a more precise prediction, applied the biomarker to represent Mr. Lund as *a case of CLL with 17p deletion*, inferring from this abstraction his prognosis. Dr. Hunt made the same inference as Dr. Roy, the resident in LEUKEMIA. But this time it was unchecked by a colleague and communicated to the patient.

As Dr. Cong, the consultant in LEUKEMIA, might have pointed out, the biomarker in this case raises more questions than answers. For one, this abstraction does away with considerable nuance and uncertainty surrounding the significance of 17p deletion in early-stage CLL. Contrary to what Dr. Hunt told his patient, it is unclear what this biomarker means for someone who is previously untreated and asymptomatic. The figures just cited, where 17p deletion reduced survival from twelve years to two, were median measures in cohorts where the majority of patients were symptomatic with more advanced, previously treated disease (Döhner et al. 2000).

This case illustrates an instance of biomarker reification. Dr. Hunt treats the biomarker as a decisive prognostic indicator for Mr. Lund, rather than as an abstraction that must be interpreted within the broader clinical context. As a result, he gives the biomarker undue salience in his clinical reasoning. The epistemic harms that follow are akin to the types discussed by Reynolds (2020). We see *epistemic capture* in how he treats the biomarker's prognostic significance as definite, despite considerable uncertainty. We also see *value partitioning* in how he classifies Mr. Lund using an overly simplistic evaluative scale, judging 'bad disease' on the basis of the biomarker alone. These can be understood as components of biomarker reification. Like the diabetes doctor who rests their entire assessment on a patient's hemoglobin A1c—inferring from its elevation 'unhealthy lifestyle' or mistaking it for a measure of health or wellbeing—the hematologist here has reified the genomic biomarker. In so doing, he violated the second maxim of clinical abstraction.

How would clinical abstraction handle this case differently? By questioning the omissions required to generate the abstraction, clinical abstraction challenges the assumptions which contributed to epistemic capture and value partitioning, thereby guarding against reification. To begin with, it questions the primacy of representing Mr. Lund as *a case of CLL with 17p deletion*, where this representation is given independent standing and applied to the exclusion of other interpretive resources. Such resources might draw attention to other features of this case, say, the clinical evidence of indolent disease or Mr. Lund's own experiences. Clinical abstraction would lead Dr. Hunt to consider other possible representations of his patient. He might, for example, see Mr. Lund as an individual with early-stage, previously untreated disease who was underrepresented in the original cohorts on which he based his earlier prognosis (Döhner et al. 2000). This, in turn, might guide how Dr. Hunt discusses prognosis with Mr. Lund, doing so with an epistemic humility that acknowledges the uncertainties inherent to his case.

In short, clinical abstraction tells us that Mr. Lund should only be labelled as *a case of CLL with 17p deletion* insofar as this abstraction serves to resolve uncertainty and guide action in the case at hand. In this case, it did not serve this function. The omissions required to move from the patient to the abstraction rendered the abstraction otiose, ill-suited to resolve uncertainty and guide prognosis. Its use in this way violated the first maxim of clinical abstraction.

It is, of course, possible that the abstraction *CLL with 17p deletion* could, in some cases, help resolve uncertainty and guide action. Thinking about a patient as *a case of CLL with 17p deletion* can indeed prove useful in the right circumstances. Even in this case, it might have helped resolve uncertainty if, say, Dr. Hunt were equipped with cohorts that more closely resembled Mr. Lund, limited his inferences to those cohorts, and communicated accordingly.

Nevertheless, other features of genomic biomarkers make them especially prone to reification, even when used for ostensibly narrower purposes and more limited inferences. This is because of the multiple steps required to generate an abstract representation using genomic biomarkers, steps which are liable to be overlooked in practice. Looking more closely at the abstraction *CLL with 17p deletion*, there is considerable clinical heterogeneity in patients named by this abstraction. This is owing to several factors; not only, as we have seen, whether the biomarker is identified at diagnosis or acquired following treatment, but also

the size of the leukemic clone bearing the 17p deletion and the genomic context in which it occurs.

These latter two factors in particular highlight the complexities interpreting genomic biomarkers in practice. Firstly, the size of the leukemic clone harbouring the 17p deletion is known to be an independent predictor of clinical outcomes (Delgado et al. 2012). Despite this, the variable is often reported and interpreted by clinicians in a binary fashion—*positive vs. negative*—although cut-offs for positivity for this biomarker are not standardized across laboratories or studies.[25]

Secondly, the significance of the 17p deletion depends not only on the clinical context, but also on genomic context, with concomitant genetic mutations impacting survival (Yu et al. 2017). This context dependence is an increasingly recognized feature of genomic biomarkers, which problematizes prediction on the basis of a single biomarker and complicates prognostication. This has led to increasingly complex systems for genomic risk stratification, now seen across a number of cancers.[26]

Precision oncology aims to overcome this prognostic underdetermination with research built on large-scale sequencing efforts. However, in the case of CLL, advances in molecular diagnostics enabling direct sequencing of the *TP53* gene further complicated matters, leading to widespread conflation of the 17p deletion and *TP53* mutations in guidelines and prognostic scoring systems, despite important distinctions between these categories of biomarkers.[27] These factors highlight just a few of the omissions required to move from a patient in the clinic to their representation in terms of a genomic biomarker. All of this complexity was abstracted away when Dr. Hunt called Mr. Lund *a case of CLL with 17p deletion*.

Some might see these challenges as inevitable pitfalls in knowledge translation between basic science and clinical practice. Others might argue that such detailed attention to the nuances of biomarker measurement is impractical for busy clinicians, who necessarily rely on crude categories and imperfect abstractions for the purposes of clinical expediency. Maxim three, they might argue, is simply too demanding. Doctors can't spend their days 'questioning omissions'—there are patients to see!

As should be clear by now, my objection is *not* to the use of abstractions as practical, expedient tools to facilitate clinical reasoning, communication, and action. Indeed, clinical abstraction stresses abstraction's practical utility. Rather, my objection is to treating abstractions—qua tools for reasoning—as something more, giving them standing and scope which forgets their grounding in the concrete situation. In medicine, though expediency is valued, due to reification's prudential and epistemic harms, abstraction must be especially judicious. Put more strongly, clinicians have an epistemic and ethical responsibility to avoid reification. Clinical abstraction, therefore, is integral to responsible clinical reasoning.

---

[25] This is an example of a line-drawing problem arising with biomarkers (c.f., Schwartz 2007).

[26] The Genetically Inspired Prognostic Scoring System (GIPSS) in myelofibrosis, discussed in §2.4, is one such example.

[27] See, for discussion, Chin-Yee et al. (2020). Moreover, similar line-drawing problems arise with *TP53* mutations, for which the clinical significance of minute, 'subclonal' mutations remains to elucidated (Nadeu et al. 2016).

Practicing clinical abstraction, to be sure, isn't easy. Its maxims are indeed demanding, prone to violation by novices and experts alike. It's nevertheless a core competency in medicine, painstakingly learned through years of training. It's an approach that the consultant, Dr. Cong, cultivated over countless hours, and that her resident, Dr. Roy, will learn to replicate with regular practice. In this way, clinical abstraction offers both a descriptive and normative account, outlining key epistemic and ethical responsibilities in clinical reasoning while also describing what the best doctors already do.

4 Conclusion

The growing use of biomarkers in precision medicine and oncology raises concerns over potential epistemic harms posed by genomic technologies. This article has examined the role of biomarkers in clinical reasoning, focusing on the epistemic harms of reification and strategies to mitigate them. I proposed *clinical abstraction* as a descriptive and normative account of reasoning with biomarkers, which captures how clinicians use biomarkers as tools to guide decision-making under uncertainty while also maintaining sensitivity to the harms of reification. While this article has focused on the uses and abuses of biomarkers qua abstractions, it opens the door to further research on the role of abstraction in clinical reasoning more generally, with implications for philosophers, educators, and clinicians alike. Moreover, this article illustrates how precision medicine and oncology can benefit from interdisciplinary engagement between philosophers and clinicians to help address emerging epistemic and ethical issues in patient care.

Acknowledgments

Thanks to Jacob Stegenga, Henrik Røed Sherling, Ina Jäntgen, Helene Scott-Fordsman, Cristian Larroulet Philippi, Stephen John, Anya Plutynski, Jonathan Fuller, and two anonymous reviewers for helpful feedback. Thanks also to audiences at the International Philosophy of Medicine Roundtable, British Society for the Philosophy of Science Annual Conference, and the Philosophy and Methodology of Medicine Workshop at the Munich Centre for Mathematical Philosophy for thoughtful comments and discussion.

References

Blanchard, A. 2016. Mapping Ethical and Social aspects of Cancer Biomarkers. *New Biotechnology* 33: 763-772.

Boyce, A. 2019. "Genomic Contextualism, Genetic Determinism, and Causal Models." *American Journal of Bioethics* 19 (1): 73-75.

Callahan K.P., Flibotte J., Skraban C. et al. 2022. Influence of Genetic Information on Neonatologists' Decisions: A Psychological Experiment. *Pediatrics* 149 (3): e2021052130.

Cartwright, N. 1994. *Nature's Capacities and Their Measurement*. Oxford: Oxford University Press.

Cartwright, N., and H. Mendell. 1984. "What Makes Physics' Objects Abstract?" In *Science and Reality*, edited by J.T. Cushing, C.F. Delaney, and G.M. Gutting, 134-152. Notre Dame: University of Notre Dame Press.

Cervantes, F., et al. 2009. "New Prognostic Scoring System for Primary Myelofibrosis Based on a Study of the International Working Group for Myelofibrosis Research and Treatment." *Blood* 113 (13): 2895-2901.

Chin-Yee, B., B. Sadikovic, and I.H. Chin-Yee. 2020. "Genomic Data in Prognostic Models—What Is Lost in Translation? The Case of Deletion 17p and Mutant TP53 in Chronic Lymphocytic Leukaemia." *British Journal of Haematology* 188 (5): 652-660.

Chin-Yee, B., and A. Plutynski. 2023. "Concepts of Actionability in Precision Oncology." *Philosophy of Science* 1-12.

Chin-Yee, B. 2024. "Cancer Medicine and Precision Oncology." In *Handbook of the Philosophy of Medicine*, edited by M. Walker and T. Schramme. Springer.

Collins, F.S., and A.D. Barker. 2007. "Mapping the Cancer Genome." *Scientific American* 296 (3): 50-57.

Croskerry, P. 2002. Achieving Quality in Clinical Decision Making: Cognitive Strategies and Detection of Bias. *Academic Emergency Medicine* 9: 1184-1204.

Custers, E.J. 2015. "Thirty Years of Illness Scripts: Theoretical Origins and Practical Applications." *Medical Teacher* 37 (5): 457-462.

Dam, M.S., S. Green, I. Bogicevic, et al. 2022 Precision Patients: Selection Practices and Moral Pathfinding in Experimental Oncology. *Sociology of Health & Illness* 44:345-359.

Delgado, J., et al. 2012. "Chronic Lymphocytic Leukaemia with 17p Deletion: A Retrospective Analysis of Prognostic Factors and Therapy Results." *British Journal of Haematology* 157 (1): 67-74.

Döhner, H., et al. 2000. "Genomic Aberrations and Survival in Chronic Lymphocytic Leukemia." *New England Journal of Medicine* 343 (26): 1910-1916.

FDA-NIH Biomarker Working Group. 2016. *BEST (Biomarkers, EndpointS, and Other Tools) Resource*. Washington, DC: Food and Drug Administration and National Institutes of Health.

Garrison, N.A., et al. 2019. "Genomic Contextualism: Shifting the Rhetoric of Genetic Exceptionalism." *American Journal of Bioethics* 19 (1): 51-63.

Godfrey-Smith, P. 2009. "Abstractions, Idealizations, and Evolutionary Biology." In *Mapping the Future of Biology*, edited by A. Barberousse, M. Morange, and T. Pradeu, 47-56. Springer.

Gould, S.J. 2006. *The Mismeasure of Man*. New York: W. W. Norton & Company.

Hájek, Alan. 2007. "The Reference Class Problem is Your Problem Too." *Synthese* 156 (3): 563–85.

Hey, S.P. 2015. "Judging Quality and Coordination in Biomarker Diagnostic Development." *THEORIA: Revista de Teoría, Historia y Fundamentos de la Ciencia* 30 (2): 207-227.

Hey, S.P., and B. Barsanti-Innes. 2016. "Epistemology, Ethics, and Progress in Precision Medicine." *Perspectives in Biology and Medicine* 59 (3): 293-310.

Ishak, N.D., et al. 2023. "Cancer Patients' Experience of Receiving Variant of Uncertain Significance Results: An Asian Perspective." *Journal of Genetic Counselling* 1-10.

Jones, M.R. 2005. "Idealization and Abstraction: A Framework." In *Idealization XII: Correcting the Model*, edited by M.R. Jones and N. Cartwright, 173-217. Amsterdam: Rodopi.

Kaplan, J.M., and R.G. Winther. 2013. "Prisoners of Abstraction? The Theory and Measure of Genetic Variation, and the Very Concept of 'Race'." *Biological Theory* 7: 401-412.

Kerr, A., J. Swallow, C.K. Chekar, S. Cunningham-Burley. 2019. Genomic Research and the Cancer Clinic: Uncertainty and Expectations in Professional Accounts. New Genetics and Society. 38 (2): 222-239.

Lane, A.S., K.E. Lynch, M. Arnold et al. 2023. The Undue Influence of Genetic Information on Senior Medical Students' Treatment Decisions. *BMC Medical Education*. 23:938.

Leptak, C., et al. 2017. "What Evidence Do We Need for Biomarker Qualification." *Science Translational Medicine* 9 (417): 1-4.

Levy, A. 2018. "Idealization and Abstraction: Refining the Distinction." *Synthese* 198(24): 5855-5872.

Li, M.M., et al. 2017. "Standards and Guidelines for the Interpretation and Reporting of Sequence Variants in Cancer: A Joint Consensus Recommendation of the Association for Molecular Pathology, American Society of Clinical Oncology, and College of American Pathologists." *Journal of Molecular Diagnostics* 19 (1): 4-23.

Macklin S.K., Jackson J.L., Atwal P.S., Hines S.L. 2019. Physician Interpretation of Variants of Uncertain Significance. *Familial Cancer* 18: 121-126.

Makhnoon, S., B.H. Shirts, and D.J. Bowen. 2019. "Patients' Perspectives of Variants of Uncertain Significance and Strategies for Uncertainty Management." *Journal of Genetic Counseling* 28 (2): 313-325.

Makhnoon, S., et al. 2021. "Perceptions of Provider's Epistemic Authority in Response to Variant of Uncertain Significance-Related Recommendations." *Journal of Genetic Counseling* 30: 513-521.

Nabhan, C., R. Gordana, Y.L. Wang. 2015. Predicting Prognosis in Chronic Lymphocytic Leukemia in the Contemporary Era. *JAMA Oncology*. 1(7): 965-974.

Nadeu, F., et al. 2016. "Clinical Impact of Clonal and Subclonal *TP53*, *SF3B1*, *BIRC3*, *NOTCH1*, and *ATM* Mutations in Chronic Lymphocytic Leukemia." *Blood* 127 (17): 2122-2130.

Plutynski, A. 2022. "Why Precision Oncology Is Not Very Precise (And Why This Should Not Surprise Us)." In *Personalized Medicine in the Making: Philosophical Perspectives from Biology to Healthcare*, edited by C. Beneduce and M. Bertolaso. Cham: Springer.

Plutynski, A. Forthcoming. "The Cancer Genome Atlas Project: Natural History, Experiment, or Something In-Between?" In *Perspectives on the Human Genome Project and Genomics*, edited by C. Donohue and A.C. Love. Minneapolis: University of Minnesota Press.

Potochnik, A. 2017. *Idealization and the Aims of Science*. Chicago: University of Chicago Press.

Reichenbach, H. 1949. "Philosophical Foundations of Probability." *Proceedings of the Berkeley Symposium on Mathematical Statistics and Probability*. University of California Press: 1–21.

Reynolds, J.M. 2020. "What If There's Something Wrong with Her? How Biomedical Technologies Contribute to Epistemic Injustice in Healthcare." *The Southern Journal of Philosophy* 58 (1): 161-185.

Salmon, W.C. 1971. *Statistical Explanation and Statistical Relevance*. Pittsburgh: University of Pittsburgh Press.

Scherr, C.L., et al. 2021. "How Patients Deal with an Ambiguous Medical Test: Decision-Making After Genetic Testing." *Patient Education and Counseling* 104 (5): 953-959.

Schwartz, P.H. 2007. "Defining Dysfunction: Natural Selection, Design, and Drawing a Line." *Philosophy of Science* 74: 364-385.

Spivak, J.L. 2017. "Myeloproliferative Neoplasms." *New England Journal of Medicine* 376 (22): 2168-2181.

Strevens, M. 2008. *Depth: An Account of Scientific Explanation*. Cambridge: Harvard University Press.

Tabery, James. 2024. *Tyranny of the Gene: Personalized Medicine and its Threat to Public Health*. New York: Alfred A. Knopf.

Tabb, K., and M. Lemoine. 2021. "The Prospects of Precision Psychiatry." *Theoretical Medicine and Bioethics* 42 (5-6): 193-210.

Tefferi, A., et al. 2018. "GIPSS: Genetically Inspired Prognostic Scoring System for Primary Myelofibrosis." *Leukemia* 32 (7): 1631-1642.

Thomas, J.M., L.M. Cooney Jr., and T.R. Fried. 2019. "Prognosis Reconsidered in Light of Ancient Insights—From Hippocrates to Modern Medicine." *JAMA Internal Medicine* 179 (6): 820-824.

Timmermans, S., C. Tietbohl, E. Skaperdas. 2017. Narrating Uncertainty: Variants of Uncertain significance (VUS) in Clinical Exome Sequencing. *BioSocieties* 12: 439-458.

Tonelli, M.R., and B.H. Shirts. 2017. "Knowledge for Precision Medicine: Mechanistic Reasoning and Methodological Pluralism" *Journal of the American Medical Association* 318 (17): 1649-1652.

Tversky, A., and D. Kahneman. 1974. Judgment under Uncertainty: Heuristics and Biases. *Science* 185: 1124-1131.

Walsh, D.M. 2020. "Action, Program, Metaphor." *Interdisciplinary Science Reviews* 45 (3): 344-359.

Wasserman, D.T., and A. Asch. 2005. Where is the Sin in Synecdoche?: Prenatal Testing and the Parent-Child Relationship. In *Quality of Life and Human Difference: Genetic Testing, Health Care, and Disability*, edited by D. Wasserman, J. Bickenbach, and R. Wachbroit. Cambridge: Cambridge University Press.

Winther, R.G. 2014. "James and Dewey on Abstraction." *The Pluralist* 9 (2): 1-28.

Winther, R.G. 2020. *When Maps Become the World*. Chicago: University of Chicago Press.

Woyach, J.A., et al. 2018. "Ibrutinib Regimens versus Chemoimmunotherapy in Older Patients with Untreated CLL." *New England Journal of Medicine* 379 (26): 2517-2528.

Yu, L., et al. 2017. "Survival of del17p CLL Depends on Genomic Complexity and Somatic Mutation." *Clinical Cancer Research* 23 (3): 735-745.

Zhong, L., E.E. Donovan, and A.L. Vangelisti. 2021. "Examining the Effectiveness of Genetic Counselors' Communication of Variant of Uncertain Significance Results of Breast Cancer Genes." *Health Communication* 36 (5): 606-615.