

**Epistemological and Ethical Aspects of Time
in Scientific Research**

Von der Philosophischen Fakultät
der Gottfried Wilhelm Leibniz Universität Hannover
zur Erlangung des Grades Doktor der Philosophie

Dr. phil.

genehmigte Dissertation von

Daria Jadreškić

Erscheinungsjahr 2020

Referent: Prof. Dr. Torsten Wilholt

Institut für Philosophie, Leibniz Universität Hannover

Korreferentin: Prof. Dr. Anke Büter

Department of Philosophy and History of Ideas, Aarhus University

Tag der Promotion: 10.06.2020

Acknowledgments

I would like to thank my supervisors Torsten Wilholt and Anke Büter for guiding me through the dissertation process in the most academically rigorous, yet most possibly encouraging manner. Conversations with Torsten Wilholt have especially contributed to my argumentation skills, broadness of interest, attention to style, reader-friendliness, and most of all, confidence that this dissertation is achievable. Anke Büter has been a careful reader and an attentive listener of my talks at many conferences that we have attended together. Her feedback has made my work better in various ways every time she provided it and I have learned a lot from being a regular attendant of her talks as well.

I am grateful to my research training group DFG-GRK 2073 Hannover-Bielefeld for letting me be a part of the team. I think we are a great crowd of people, who have always managed to make our common research business into a friendly team-building process that has brought me so much fun and insight. I could not have picked a better place to get this thing done. Our research colloquia, theory workshops, and peer group meetings are events that I have always enjoyed, for the talks as well as for the dinners.

The Hannover part of the group has basically been my second home, next to my home in Croatia. I am quite sure that I have spent most of the time in the office these three years, not necessarily in order to work, but in order to hang out with friends. I cherish very fond memories from both of our addresses. I would like to thank Rafael Ugarte Chacón for his patience with my questions and attention to various administrative details that I would have never been able to catch up with on my own. I also thank him for the reliable comfort of the morning office coffees which have been a much-needed constant in the ever-changing world. I have learned so much about so many things there. My special thanks go to David Hopf and Stefano Canali for being loyal members of our coffee group. Both of them have been an especially invaluable resource for topics, philosophical exchange, and extreme human kindness. Many thanks to Julia Pfeiff for being a perfect office mate and a dear friend. I don't think I have ever loved any disagreements more than ours. Thanks to Enno Fischer for his friendly support and for being a model in dedication to philosophy, but also in work-life balance. His young family has been a breath of fresh family air that has been a pleasure to see growing.

My gratitude also goes to other members of the Hannover group and the Institute of Philosophy who have made my doctoral studies richer: Anna Leuschner, Philipp Eichenberger, Leon Schäfer, Tobias Schönwitz, Markus Ahlers, Jonas Lipski, Lukas Steinbrink, and Simon Lohse, with a special credit to Martin Wasmer and Lucie White for their readiness to join in for a beer when needed. I am also grateful for the privilege to get to know the Carson-Feest family that has made my Hannover life much warmer. Thanks to Mathias Frisch, Ralf Stöcker, and Saana Jukola for their valuable feedback, as well as to other members of the graduate school who have on various occasions had the opportunity to engage with my work.

The Institute of Philosophy in Hannover has provided me the most motivating academic environment. Thanks to students in my seminar for the fruitful discussions we had. Thanks to Jessica Inhestern and Güliz Erdem for helping me out with organizational and administrative issues.

My special thanks go to three researchers from the Albert Einstein Institute in Hannover who have been kind enough to introduce me to the fascinating world of gravitational wave physics and answer my questions about their career and work, as well as to their colleagues who have guided me through the GEO600, a gravitational wave detector near Hannover, during the Open Day on June 16, 2019.

I am indebted to the audiences at many conferences I had the opportunity to attend. I have been especially encouraged and have benefited from the feedback by and conversations with Jeffrey Kenneth Aronson, Philip Kitcher, Kevin C. Elliott, Barbara Osimani, Miriam Solomon, Rebecca Kukla, Rachel Ankeny, Mark Tonelli, Chrysostomos Mantzavinos, Paul Hoyningen-Huene, and Bennett Holman. My special thanks go to Tomasz Rzepiński for inviting me to give a talk at the Adam Mickiewicz University in Poznań and for his generous hospitality during my stay there, as well as to Marko Jurjako for inviting me to give a talk at my hometown University of Rijeka.

I am grateful to dear friends I have made in Hannover, who have been an invaluable support in many aspects of my life here: Radhika Natarajan and Siham Schotemeier. I also thank many others with whom I have crossed paths and established a connection that I will always cherish: Amélie Fau, Sahar Dehnavi, Iris Phan, Vedat Uner, and Edith Velázquez. I have been especially indebted by the enormous hospitality of the Bošnjak family and the many dear ladies and gentlemen working in the Conti-Campus Contine.

This is as far as Germany goes, and my acknowledgements now travel southeast, to Croatia. My connection with Hannover would have never been established had it not been for the project designed and executed at the University of Rijeka by Predrag Šustar, Smilja Cukrov, and Zdenka Brzović, and later led by Elvio Baccarini. I am indebted to all of them for involving me in the team, as well as to Nenad Smokrović and Snježana Prijić-Samaržija for always being supportive of my academic work. Special thanks in this respect go to Bruno Lončarić and s. Veronika Mila Popić to whom I am grateful for their generosity, openness, and active support during many years of our friendship that has outgrown our professional connections.

I am indebted to many dear friends who have always supported me: Marija Sekulić, Marijana Dlačić, Diana Gašpić, Ivana and Srđan Baričević, Silvia Vlaše, Blanka Krivičić, Jan Hyrat, Nataša Jakšić, Jasmina Kovačić, Iva and Ivan Bujger, Martina Blečić, Vedran Krušvar, Boris Ružić, Saša Stanić, Stella Defranza, Viktorija Žgela, Dinko Benčić, Igor Eterović, and the Veljak family.

Finally, my greatest “Thank you” goes to my family: mother Dubravka, father Dario, sister Ornella, and brother Jakov. They have been my strength and motivation. The love and support I have always had from them, as well as from my grandparents, is beyond words. Hvala vam na svemu.

Abstract

This dissertation explores the influence of time constraints on different research practices. The first two parts present case studies, which serve as a basis for discussing the epistemological and ethical implications of temporal limitations in scientific research. Part I is a case study on gravitational wave research, conducted by the LIGO Scientific Collaboration. This exemplifies fundamental research – without immediate societal applications, open-ended in terms of timeline and in terms of research goals. It is based, in part, on qualitative interviews conducted with gravitational wave physicists. I show that considerations about time and speed play a role in every stage of research: goal setting, method design, and the evaluation and communication of results. Part II provides a case study on translational medicine, an approach explicitly dedicated to accelerating research in order to develop and implement new therapies. This epitomizes applied research with high social stakes, motivated by non-epistemic goals. Here, epistemic trade-offs between speed and reliability intersect with ethical trade-offs between different types of harms. In Part III, the insights from both of these case studies are used as the basis for a more general discussion concerning the pragmatic aspects of epistemic practices, especially in relation to current debates centered on the role of values in science. A particular focus is on the value of speed and the ability to generate reliable results, either via choice of methods, or via decisions about which goals to set, as well as decisions about when to stop further testing. The primary thesis of the dissertation is that pragmatic considerations stemming from limitations of resources are a necessary feature of the pursuit of epistemic aims, and that the epistemic is thus inherently pragmatic.

Keywords: scientific practice; gravitational wave research; translational medicine; values in science; speed; reliability; ethical implications of research

Contents

Introduction

1. An unfair comparison.....1
2. The problem.....4
3. Overview of the dissertation.....7

Part I

Case Study I: Gravitational Wave Physics

1. Introduction: LIGO.....10
2. Speed of research and goal setting.....21
3. Prioritization and urgency.....31
4. Time and methods.....38
 4.1 Matched filtering.....38
 4.2 Machine learning.....43
 4.3 Searches for continuous waves.....45
5. Time and the evaluation and communication of results.....48
 5.1 Introduction.....48
 5.2 Non-epistemic and pragmatic values in the communication of results.....50
6. Conclusion.....56

Part II

Case Study II: Translational Medicine

1. Introduction.....59
2. What is translational medicine? 63
 2.1 The location and scope of translational gap(s).....65
 2.2 Three models of translational process.....67
 2.3 Causes of translational gap(s).....73
3. A short history of translational medicine.....76
4. Accelerating discovery and research – hopes and tensions.....82

5. The landscape of biomedical research.....	88
5.1 Personalized or precision medicine.....	88
5.1.1 The relationship between personalized medicine and translational medicine.....	91
5.2 Evidence-based medicine.....	95
5.2.1 The relationship between evidence-based medicine and translational medicine.....	97
5.3 Conclusion.....	105
6. Limitations to the acceleration of biomedical research.....	107
6.1 The case of early drug discovery.....	107
6.2 Measuring the effectiveness of medical interventions.....	113
7. Ethical considerations related to translational medicine.....	116
8. Research on cortisone in the 1930s-1950s.....	122
8.1 Introduction.....	123
8.2 Discovery and translation.....	124
8.3 Implementation and problems.....	128
8.4 Discussion.....	130
9. Conclusion.....	135

Part III

Time, Science, and Values

1. Introduction.....	143
2. Epistemic and non-epistemic values.....	147
3. Pragmatic values.....	153
3.1 Ernan McMullin’s account.....	153
3.2 Heather Douglas’ account.....	156
3.3 Matthew Brown’s account.....	157
4. Time-sensitivity in science.....	160
4.1 Introduction.....	160
4.2 Ease of use contributes to speed of generating results.....	163
4.3 Speed as an epistemic value.....	165
4.4 Extrinsicly epistemic is equal to non-epistemic-but-legitimately-influencing.....	167

4.5 What is a value?177

4.6 Time-sensitivity and transient underdetermination172

5. Time-sensitivity and the two case studies.....174

6. Conclusion.....177

Conclusion.....179

Bibliography.....181

Introduction

1. An unfair comparison

In 2015, two important scientific events took place within only a few months. One was the phase III trial for testing recombinant vesicular stomatitis virus-Zaire Ebola vaccine (rVSV-ZEBOV) against Ebola virus disease. The trial, called Ebola ça Suffit! (Ebola, that's enough!), began in Guinea in March, and the results suggested that the experimental vaccine could protect immunized individuals after a delay of about 10 days (Henao-Restrepo et al. 2015, 2017). The randomization stopped on July 26, after only four months, to allow for more adults to receive the vaccine immediately and to include younger age groups sooner (WHO 2015). The vaccine was approved for “compassionate use” in outbreaks, meaning that it had been proven sufficiently safe and effective to be recommended, though it had not yet been formally approved by a full regulatory process (Calain 2018). According to recent correspondence in the *Lancet*, the efficacy estimate of the vaccine remains at 100% despite doubts about bias in the research design (Henao-Restrepo et al. 2018 to Metzger and Vivas-Martínez 2018). The vaccine eventually contributed to controlling the 2013-2016 Ebola virus disease epidemic in Guinea (Geisbert 2017; Calain 2018) and it has been used in recent efforts to control the disease outbreak in the Democratic Republic of the Congo (WHO 2018).

The other event was the first direct detection of gravitational waves, on September 14 at 5:51 a.m. Eastern Daylight Time by both of the two Laser Interferometer Gravitational wave Observatory (LIGO) detectors, located in Livingston, Louisiana, and Hanford, Washington, USA (Abbott et al. 2016a; LIGO 2016). The detection was not announced until February 2016, after it had been determined that the event had actually occurred.¹ The detected gravitational waves were produced 1.3 billion years ago by the merging of two black holes into a single, more massive black hole. The detection confirmed a prediction of Albert Einstein's 1915 general theory of relativity, and three scientists involved in the discovery won the Nobel Prize for physics the next year. By September 2019 ten binary black hole mergers and one neutron

¹ For a sociological perspective on the first detection, its analysis, announcement, and the process of writing the detection paper see Collins (2017).

star merger have been observed and announced by the LIGO Collaboration and the Virgo (the European observatory for detecting gravitational waves).

Ebola ça Suffit! was a result of collective efforts to respond to the 2013-2016 West African Ebola epidemic that had caused the death of more than 11.000 people (Calain 2018, 3). In August 2014 the Ebola epidemic was declared a public health emergency of international concern and the World Health Organization (WHO) set up a panel of experts to consider the ethical permissibility of testing potentially effective interventions for the disease in an accelerated manner. Within a few months, a number of novel or repurposed therapeutic agents were tested for efficacy at various locations experiencing an outbreak. Phase II/III trials of Ebola vaccines were also conducted (Calain 2018, 3). The design of these trials was not standard, due to time constraints, limited amount of vaccine supplies, ethical concerns regarding the adoption of research methodology, and logistics and field operational challenges (Varghese 2018; Calain 2018). The Ebola ça Suffit! ring trial used cluster randomization over individually controlled randomization and a delayed vaccination arm as the control group over placebo control group, in order to maximally mitigate the transmission of the disease. Upon confirming a case of the Ebola virus, a ring (cluster) of all the infected person's contacts was established, as well as the contacts of their contacts. The clusters were assigned to either an immediate vaccination or delayed vaccination arm of the trial (for all eligible individuals aged 18 years or more and not pregnant, breastfeeding, or severely ill). The delayed arm was treated after 21 days, which allowed for both groups to receive the vaccine, as opposed to treating the control group with a placebo (Henao-Restrepo et al. 2017, 505). It was a case of research with a strong social commitment and pressing ethical considerations on the level of research design (see especially Varghese 2018), quickly adapted to the urgency of the situation and "challenging the usual benchmarks of therapeutic development" (Calain 2018, 3), because it normally takes more than a decade to develop a therapeutic intervention.

The first gravitational wave detection was a huge scientific success, opening the door to a better understanding of the universe. Gravitational waves are "ripples in the fabric of spacetime" (LIGO 2016, 1) which carry information about their origin and observing them allows scientists to "listen to the universe" (Wu 2017): learn about the nature of black holes, neutron stars, gravity, and the history of the universe. In the words of David H. Reitze, executive director of the LIGO Laboratory, reported in the news release announcing the detection:

“Our observation of gravitational waves accomplishes an ambitious goal set out over five decades ago to directly detect this elusive phenomenon and better understand the universe, and, fittingly, fulfils Einstein’s legacy on the 100th anniversary of his general theory of relativity.” (LIGO 2016, 2)

The two events are incomparable in many respects: what holds for emergencies does not hold for regular contexts, and the social commitment to improve people’s lives that is constitutive of medicine is not constitutive of physics and astronomy. I will here use the basic-applied research distinction, thinking of it in terms of a continuum, where gravitational waves research is clearly on the basic side and Ebola ça Suffit! is on the applied. I will not attempt to define basic and applied research, i.e. whether the distinction is conceived in terms of goals, motivations, accountability, or products of the research (cf. Pielke 2007, 80-96), but will rather regard all these elements as contributing to a distinction that can be used for mapping the field, despite its imprecision.² According to this “conflated” distinction, basic science is driven by curiosity, aims at understanding phenomena, produces knowledge, and is accountable to scientific peers, while applied science is driven by the desire to apply the existing knowledge to new contexts, aims at intervention and problem solving, produces technology, and is accountable to society at large.

This picture is clearly problematic when we think of our unfair comparison along these lines. Although the Ebola vaccine trial was set up primarily with the aim of assessing the efficacy of a possibly lifesaving treatment, important epistemic and ethical understanding can be expected to emerge from it. For example, understanding related to the identification of the immunological mechanism of the vaccine, or the assessment of reliability of ring design trials and, more generally, adaptive design trials. Furthermore, understanding can be broadened by attempts to specify “exceptional circumstances” and the meaning of “compassionate use”, as well as by attempts to identify the goals of interventional research in terms of individual as opposed to collective interests (see Calain 2018). It can be objected that understanding ethical challenges does not fall under the domain of scientific understanding, but in this case the two are strongly intertwined. On the gravitational physics side, laser interferometry has been applied to a new context when engaging in the search for gravitational waves, and a huge technological and

² A similar usage to which a conflated version of the “context of discovery” and “context of justification” distinction is often put (see Hoyningen-Huene 2006).

engineering enterprise has contributed to this knowledge-seeking endeavor, with spin-offs being applied in fields like optics and measurement science.

Be that as it may, the unfair comparison serves to illustrate that science as a whole comprises different projects operating in different timelines out of different motivations and with different goals. We have accelerated, emergency-solving, society-serving scientific actions, as well as slower, open-ended, long-term knowledge-seeking research projects. This thesis discusses the ways in which diverse types of scientific research deal with different time constraints, and how these time constraints are reflected in goal setting, method design, and the evaluation and communication of results. Furthermore, the thesis explores the relation between time and values internal and external to science, as well as the ethical implications of the methodological choices made.

2. The problem

Whether the unit be an individual scientist, a research program, a project, a paradigm, or a scientific community of a certain generation, the predicament is this: we acquire knowledge using the time we have, and not the time we would ideally need. Each sequence of time on the axis of human history is a constraint on the availability of evidence. Time constraints can be deadlines posed by funders, self-estimated milestones, particular time slots in which an event of interest occurs and can be measured or analyzed, or an individual lifetime in which a particular epistemic goal is pursued. Hence, knowledge is acquired in a limited time frame, building on past achievements, but having a range of a lifetime (of a person, a group, a program, a paradigm, a funding cycle) to say something more, whether it is a positive, additive ‘more’, or a negative one, a falsification of what was previously considered to be true. Both the positive and the negative step represent progress because both bring the next generation closer to a more accurate description of observable phenomena, and hopefully to a better life for all. We expect these two features, accuracy and a good life, not to lead us in different directions. The predicament of “using the time we have” becomes a real problem only if available evidence suggests that time is precisely what we do or may not have, at least not without significant losses. These losses can be purely epistemological, but they can also be, and often are, social and ethical, as was the case with the Ebola disease epidemic that motivated the accelerated testing of experimental treatments. The priority in these cases is to keep the losses to a minimum by developing methods that can contribute to expedited desired outcomes.

Specific scientific problems and the time frames for dealing with them vary significantly from discipline to discipline and case to case, as we can see from the unfair comparison. Problems broader than scientific ones – be they social, societal, or environmental, present urgencies which require methods that differ from the ones we would ideally use. Ideally means unlimited with respect to the availability of resources: time, money, scientists, computational capacities, technology, material, or political will. In areas such as biomedicine, climate science, or ecology, there is an extra scientific, societal pressure to acquire knowledge which can be implemented, and which can inform actions that need to be taken more or less urgently in order to prevent detrimental consequences on the local or the global level. In biomedicine, such actions include the design of new drugs and the implementation of therapeutic practices. The war on cancer that started in the 1970s and the intensive research on Ebola and Zika viruses after the outbreak of epidemics are examples of a dedication to solve societal problems as soon as possible by means of scientific achievement for the sake of the people, societies and economies affected by them. Survival and health are not states that we can postpone.³

In climate science, predictions and projections of the Earth's future climate are being made not only for the sake of knowledge itself, but also to provide information that needs to be taken into account when deciding on governmental policies. If governments would agree on the appropriate course of action, we could hopefully affect the level of global warming in the future by intervening to reduce the level of greenhouse gas emissions now. Similar concerns have certainly motivated researchers throughout the past, but the current state of science is distinctive in its increasing need for large-scale cooperation, and the level of actual cooperation in contemporary research contexts, exemplified, for example, in the work of the

³ In the context of emerging technologies which can detect predispositions for different health issues even prenatally, it can seem misleading to consider health as something achievable since we are always on a scale of probabilities for developing a certain health problem. The so-called cascade model of health therefore emphasizes prevention instead of treatment (see Boenink 2010). This model, however, raises epistemic, ethical, and regulatory questions related to the responsibility to detect and prevent health problems. There is an increasing danger of over-medicalization and overdiagnosis for issues that will never develop into a disease or an otherwise harmful condition (see Biddle 2016). Unnecessary diagnostic procedures can thus cause severe distress and harm for patients, as well as increased burden on health systems. It is outside the scope of this thesis to examine the concepts of health and disease, but it is important to emphasize the high social stakes involved in addressing disease, illness, and disability in any of their various manifestations.

Intergovernmental Panel on Climate Change (IPCC) and the Laser Interferometer Gravitational-Wave Observatory (LIGO) Scientific Collaboration (LSC). Not only horizontal cooperation between researchers is needed, but also vertical cooperation between researchers and all members of society, as it is suggested in the Responsible Research and Innovation approach.⁴ The need for and possibility of this kind of cooperation is unprecedented in the history of scientific research.

The accountability of science to society means not only addressing societal needs, but addressing them reasonably soon. In other words, there are ethical and social reasons to prefer a quicker solution to the delayed one. There are also epistemological reasons for preferring a quicker solution. If two methods are accurate, but one is faster, the faster one will make us epistemically better off in the same amount of time. However, a quick solution does not live up to its promises if it is a bad solution, i.e. if it is epistemically flawed or erroneous. Methods of risk assessment for the toxicity of substances safeguard the society from releasing them on the market (too quickly), an effort often compromised by economic interests (see Horton 2001 and Biddle 2007).⁵ Although it is clear that speed is often incompatible with rigid risk assessment for both individuals and populations, the application of scientific knowledge in practice is still needed and wanted soon, sometimes even urgently. On the other hand, there is also scientific knowledge and technology based on it that we never want to use, for example weapons of mass destruction. Science has prolonged human life significantly, but it has also made it possible to end and extinguish it altogether. It is important to draw attention to epistemological and ethical consequences of being in a predicament of a limited time frame while exploring phenomena that extends well beyond anything we can grasp: time-wise, space-wise, cognition-wise, and in the same time addressing the needs of contemporary and future biosphere. The focus of this dissertation are constraints posed by time and the way science deals with them.

⁴ Responsible Research and Innovation (RRI) approach presupposes that researchers, citizens, policy makers, business, and non-profit organizations coordinate efforts to ensure that research and innovation respond to the needs and expectations of society, reflect its values and are responsible. It is a part of 'Science in Society' program of the Horizon 2020 Research and Innovation Strategy of the European Commission (Owen, Macnaghten, and Stilgoe 2012).

⁵ Methods of risk assessment for the toxicity of substances also safeguard the society from keeping them on the market: in the United States the burden of proof is on the state to show that substances apart from pesticides and pharmaceuticals are safe. Speed of assessment plays a role in this context too (see Cranor 1993, 1995).

3. Overview of the dissertation

The dissertation is comprised of three parts. Part I and II present case studies which serve as a basis for discussing the epistemological and ethical implications of temporal limitations in scientific research. In Part III, the insights from these case studies are taken up for a more general discussion of pragmatic considerations in epistemic practices, such as considerations about time and efficiency, especially in relation to current debates centered on the role of values in science.

The main thesis of the dissertation is that pragmatic considerations stemming from limitations of resources are a necessary feature of the pursuit of epistemic aims and that the epistemic is thus inherently pragmatic. Since this pragmatic dimension has sometimes been conflated with the issue of the appropriate role for non-epistemic values in theory assessment, the aim of the dissertation is to emphasize the limitations of resources, primarily time, as important factors for understanding decision making in science, regardless and apart from any intrusion of non-epistemic values. Particularly the value of speed, instantiated in methods and practices, has been underrepresented in discussions about values in science. A recent debate which engages with the question of the epistemic status of speed is addressed, and I argue in favor of an understanding of speed as captured by the category of epistemic values. Nonetheless, I put forward a stance which accommodates a functional role for non-epistemic values in theory assessment via their influence in determining the degree of *time-sensitivity*, which is taken to be a feature of problems in their particular contexts. The degree of time-sensitivity is influenced by non-epistemic values, and at the same time influences epistemic trade-offs between speed and reliability.

In the focus of Part I is gravitational wave research conducted by the LIGO Scientific Collaboration, which has recently succeeded in directly observing gravitational waves for the first time. It is an example of basic research, open-ended in terms of timeline, but also in terms of further goals. After introducing the LIGO Collaboration (Ch. 1. *Introduction: LIGO*), I show how considerations about the time frame in which results can be acquired influence the setting of goals (Ch. 2: *Speed of research and goal setting*). In Chapter 3, the focus turns to the prioritization of research lines and goals on different levels of the gravitational wave research community: the individual level, group level, and the level of the whole research field. In this chapter I identify epistemic and non-epistemic urgencies in gravitational wave research (Ch. 3: *Prioritization and urgency*). In Chapter 4, I discuss methods in gravitational wave research and

show how they are designed to respond to the problem of limitations in time and computing power (Ch. 4: *Time and methods*). In Chapter 5, I show how time plays a significant role in the evaluation and communication of results in gravitational wave physics (Ch. 5: *Time and the evaluation and communication of results*). Finally, Chapter 6 concludes with the gravitational wave physics case study.

The focus of Part II is translational medicine, an approach explicitly dedicated to accelerating discovery and research in order to develop and implement new therapies. After introducing this approach and its goals (Ch. 1: *Introduction*), I provide an overview of different ways of conceptualizing translational medicine: through the location and scope of translational gap(s), different models of the translational process, and the causes of perceived translational gaps (Ch. 2: *What is translational medicine?*). In Chapter 3, I give an overview of the history of translational medicine (Ch. 3: *A short history of translational medicine*) and in Chapter 4, I identify some hopes and promises, but also tensions of the approach (Ch. 4: *Accelerating discovery and research – hopes and tensions*). Since biomedical research involves diverse stakes of patients, producers, regulators, and researchers, I highlight their possible conflicting interests and values as factors that might make acceleration or the means of achieving it incompatible with the goals of all parties. In Chapter 5, I present two approaches necessary to understand the epistemic means of acceleration in contemporary biomedical context: evidence-based medicine and personalized medicine (Ch. 5: *The landscape of biomedical research*). Although it is more or less agreed upon in the literature that translational medicine does not offer much novelty in terms of epistemology (Solomon 2011, 2015; Robinson 2019), this does not mean that it does not influence research outcomes in other ways (this is especially argued for in Robinson 2019). In Chapter 5 I show why the fastest contemporary translations are most likely to happen in the domain of personalized medicine. Moreover, I argue that translational initiatives have facilitated a transition towards personalized medicine, especially given the criteria of acceleration and efficiency. In Chapter 6, some limitations to the acceleration of biomedical research on the level of drug discovery and on the level of effectiveness assessment are identified, and I argue that translational medicine is a structured effort to foster serendipitous discoveries (Ch. 6: *Limitations to the acceleration of biomedical research*). In Chapter 7, an overview of ethical discussions of translational medicine is provided, including considerations at the intersection of ethics and epistemology, such as patient involvement in translational contexts (Ch. 7: *Ethical considerations related to translational medicine*). Chapter 8 presents a historical case study of a successful translation in the pre-translational era, the research on

cortisone (Ch. 8: *Research on cortisone in the 1930s-1950s*). The case of cortisone research points to some general limitations to the goal of acceleration in the biomedical context, but also shows the differences in the broader socio-economic context in which biomedical research was undertaken in the 1930s-1950s as opposed to now. Finally, Chapter 9 gives an overall conclusion of the translational medicine case study.

Part III provides a more general discussion of time, science, and values. After the introductory chapter, Chapter 2 (*Epistemic and non-epistemic values*) and Chapter 3 (*Pragmatic values*) present the concepts of epistemic, non-epistemic, and pragmatic values, which are sometimes used in the discussion of case studies. A special focus is on pragmatic values, which is a category included in the accounts of Ernan McMullin (1982) and Heather Douglas (2013). Matthew Brown's (2013) account is also presented. Although he does not establish a category of pragmatic values, he puts forward the account of pragmatist functionalism about enquiry. This account acknowledges a legitimate role for a very diverse set of values in scientific research, which is exactly what the two case studies in this thesis purport to exemplify. In Chapter 4, I put forward and analyze the concept of *time-sensitivity*, by drawing on a discussion between Kevin Elliott and Daniel McKaughan (2014) and Daniel Steel (2016) on the role of non-epistemic values in theory assessment and the epistemic status of speed (Ch. 4: *Time-sensitivity in science*). I argue that speed is an epistemic value, and propose an account of time-sensitivity. In Chapter 5, I apply the concept of time-sensitivity to the two case studies with the help of Justin Biddle's (2013b) terminology (Ch. 5: *Time-sensitivity and the two case studies*). Chapter 6 concludes Part III, after which a final concluding chapter wraps up the main theses of the dissertation.

In Part III, it is argued that an estimated degree of time-sensitivity should be understood as a contextual factor which bears on considerations about time limits, and very directly informs the ways in which epistemic aims are pursued and achieved, namely, by influencing trade-offs which involve speed. It is also argued that the concept of "contextual factors" captures the role of time-sensitivity better than the problematic concept of "values". Finally, the concept of time-sensitivity is closely related to the argument from transient underdetermination, which supports the stance that the fulfilment of epistemic aims is inherently pragmatic, as well as that non-epistemic values have a functional role in theory assessment which should not be downplayed. This stance maintains that the influence of values is necessary for socially responsible scientific work, and that the borders between different categories of values are often hard to draw.

Part I

Case Study I: Gravitational Wave Physics

1. Introduction: LIGO⁶

Laser Interferometer Gravitational wave Observatory (LIGO) Scientific Collaboration (LSC) is an international cooperation dedicated to detecting gravitational waves by means of several kilometer long ground based interferometric detectors.⁷ The first direct detection of gravitational waves was achieved on September 14, 2015⁸ (Abbott et al. 2016a; LIGO 2016).⁹ Gravitational wave detectors use laser interferometry to measure the “ripples” in spacetime – gravitational waves caused by astronomic events such as colliding neutron stars, black holes, or bursts such as supernovae – explosions of a star. Gravitational waves were predicted by Albert Einstein in 1915 by his general theory of relativity. According to it, any object with mass warps the structure of spacetime which results in other objects moving on or orbiting along the curves caused by the warping. Gravity is the curvature of spacetime. Accelerating bodies distort spacetime so that “waves” radiate from the source like ripples in a pond. These ripples are gravitational waves that then travel through the universe at the speed of light. Since

⁶ This chapter is substantively informed by the documents, media releases, and other material from the LIGO, LSC, GEO600, and Max Planck Institute for Gravitational Physics (Albert Einstein Institute) web pages. Where necessary, particular sources are cited.

⁷ As opposed to the space-based gravitational wave detector – Laser Interferometer Space Antenna (LISA), planned for launching in 2034, and an underground detector, such as the Kamioka Gravitational Wave detector (KAGRA) in Japan, due to open in late 2019 (Castelvecchi 2019). LISA Pathfinder, a precursor mission for LISA, was launched on December 3, 2015 (NASA 2015).

⁸ An indirect confirmation of the existence of gravitational waves was achieved by Russel Hulse and Joseph Taylor who observed a pair of neutron stars (a binary), whose orbit was slowly decreasing over many years, which was consistent with the emission of gravitational waves according to Einstein’s general theory of relativity. The energy loss due to the gravitational waves was observed by Taylor and Joel Weisberg (Abbott et al. 2016a, 1; Collins 2004, 1-2)

⁹ Abbott et al. 2016a is the detection paper, published in *Physical Review Letters* on February 11, 2016 and co-authored by more than 1000 members of the LIGO and Virgo Collaboration. The press conference at which the detection was announced was held the same day, and the news release (LIGO 2016) was launched.

gravitational waves produced on Earth are too small to be detected, the focus of gravitational wave research is on astronomic objects moving with extreme accelerations. Neutron stars, sometimes called pulsars,¹⁰ are caused by stars collapsing and creating high-pressure conditions in the core which makes most of the protons and electrons combine into neutrons. Black holes are extremely dense conglomerates of rotating matter which exhibit gravitational effects of such strength that nothing can escape their gravitational pull, not even light.

Interferometers work by merging two or more sources of light to create an interference pattern, which is then measured and analyzed in order to get the information about the origin of the passing gravitational wave. LIGO interferometers are L-shaped, have mirrors at the ends of the arms to reflect light and combine light beams to get an interference pattern. A laser beam is directed into a large tunnel structure, where it gets split in half and travels down both of the arms at exactly the same time. At the end of each arm, a mirror reflects the light back to where it came from, and the two beams merge back into one. Normally, they should recombine at the same time. However, when a gravitational wave passes through the light, the ripple, i.e. the little distortion of spacetime, will slightly stretch one arm of the light beam while shortening the other, and then the other way around. This particular stretching and shortening is the detected interference pattern which marks the passing of a wave, whose amplitude, frequency, wavelength, and speed is then analyzed in order to learn about its source (Abbott et al. 2016a; Collins 2004, 515-525).¹¹

LIGO includes two gravitational wave interferometers and two university research centers. The interferometers are located in Washington (LIGO Hanford) and Louisiana (LIGO Livingston), US, and the two main research centers are located at the California Institute of Technology (Caltech) in Pasadena and the Massachusetts Institute of Technology (MIT) in Cambridge. The detectors are 3002 kilometers away, but collect data simultaneously, operating as one single observatory. LIGO is funded by the US National Science Foundation (NSF) and it was its biggest investment at the time when the initial construction was approved for funding in 1992,

¹⁰ Pulsars are rapidly rotating neutron stars which emit electromagnetic radiation and are detected in the form of pulses as the star rotates (Hobbs 2019, 24).

¹¹ I thank the organizers of the Open Day at GEO600, laser interferometer near Hannover, Germany, in June 2018, for the opportunity to visit and learn about a gravitational wave detector on site. I especially thank the scientists at the GEO600 for their thorough and accessible guidance in understanding the workings of an interferometer and the physics behind the gravitational wave detection.

with 211 million dollars total, starting with 23.6 million for the same year (Thorne 1992, 208). The approval of funding for such an “esoteric” project was an issue of controversy and opposition from both inside and outside of the physics community.¹² Until now approximately 1.1 billion dollars have been invested in the observatory construction and upgrades, operational costs, and in research awards to individual scientists (NSF 2017).

The LIGO Scientific Collaboration (LSC) carries out the research of the LIGO observatories and of the GEO600 detector located near Hannover, Germany. GEO600 is designed and operated by scientists from the Max Planck Institute for Gravitational Physics and the Leibniz Universität Hannover, in cooperation with partners from the United Kingdom. There is a fourth currently active interferometer dedicated to the search for gravitational waves called Virgo (initially an Italian-French collaboration) near Pisa in Italy, operated by the EGO (European Gravitational Observatory) Collaboration. An underground detector near the city of Hida in Japan, operated by the KAGRA Collaboration, is to be opened by the end of 2019. LIGO-India is planned to be built in the Hingoli District, Maharashtra (IndiGO 2011). LSC was founded in 1997 and is currently made up of more than 1200 scientists from over 100 institutions, and from 20 countries worldwide. It is funded by both public and private sources (LSC 2018).

Construction of LIGO's original gravitational wave detectors (in Hanford and Livingston) was completed in 1999. The initial detectors started their first observing run in 2002 and ended it in 2010, with no observations made. After four years of upgrading, the two interferometers started the Advanced LIGO observation run in 2015, and within days made the first detection of gravitational waves produced 1.3 billion years ago from two colliding black holes. LIGO interferometers are identical and consist of two 4km arms each. For comparison, GEO600 in Hannover has arms of 600m length. Increasing the arm length increases the interferometer's sensitivity to vibrations, while increasing the laser power improves the interferometer's resolution. A particular feature of GEO600 is the amplification of laser light and signal, such as “squeezing” of the light and using highly reflecting mirrors to enable “power recycling” and “signal recycling” (see GEO600: “Advanced Technologies”, n.d.).

Interferometric detectors are extremely sensitive to different sources of disturbances, both external – like ground noise from the earthquakes and traffic, but also to internal laser fluctuations (Ohme 2012, 14-15). They can measure a motion 10.000 times smaller than an atomic nucleus (LIGO: “Facts”, n.d.; LIGO 2016, 4). The sum of all noise defines the sensitivity

¹² See Mervis (1991) and Collins (2004, 489-511).

of the interferometer. In addition to being the largest, LIGO interferometers are the most sensitive instruments ever built. In their most sensitive band, 100-300 Hz, the Advanced LIGO detectors are 3 to 5 times more sensitive than initial LIGO, while at lower frequencies, below 60 Hz, the improvement is tenfold (Abbott et al. 2016a, 5). The problem is that disturbances can mask or mimic a gravitational wave signal. Because of that, at least two laser interferometers located far apart are needed in order to make it possible to detect a gravitational wave signal from a binary merger, i.e. from a colliding pair of astronomic bodies. A gravitational wave signal will occur at both places at the same time,¹³ while local vibrations will occur only at one place and can therefore be disregarded. By comparing data from both sites (cross correlation), and eventually from multiple interferometers and other observatories, identical signals can be singled out as candidates for a gravitational wave detection.

Such collisions are not the only sources of gravitational waves. The 1983 LIGO Blue Book¹⁴ categorizes four gravitational wave sources which still shape the focus of the research (Collins 2004; Meadors 2014): burst (supernova), compact binary coalescence (inspiral),¹⁵ continuous wave (pulsar), and stochastic sources. The first two sources, bursts and binary coalescences, are single, transient events – violent, explosive, short, and “loud”, and therefore candidates for earlier detections since they can be singled out by cross correlation procedures.

Compact binary coalescences are exactly the kinds of events observed so far. There are three variations of such events: black hole-black hole, neutron star-neutron star, and black hole-neutron star collisions (the only ones not observed so far). The epistemic basis for the direct detection of a gravitational wave signal from these mergers is twofold: it is experimental (material, technological, and engineered) on the one side, and it is theoretical (immaterial,

¹³ More precisely, roughly at the same time. The detector in Livingston recorded the first gravitational wave signal, GW150914, 7 milliseconds before the detector in Hanford (LIGO 2016, 1), which is consistent with the prediction of gravitational waves’ propagation at the speed of light. Precision with regards to the time of the detection should be emphasized due to the importance that the considerations about simultaneity will have in this dissertation. GW150914 stands for Gravitational Wave-year (2015)-month (09)-day (14), which is the standard way of naming gravitational wave signals.

¹⁴ The 1983 Blue Book is one of LIGO’s central documents. Its real name is *A Study of a Long Baseline Gravitational Wave Antenna System*, and is written by Paul Linsay, Peter Saulson, Rai Weiss, and Stan Whitcomb in October 1983, but never published.

¹⁵ Inspiral is the first stage of a collision, in which the astronomic bodies are rotating at high speed and their orbits gradually decrease before merging.

propositional, and mathematical) on the other. It takes a properly functioning interferometer at the right sensitivity level to detect the signal, which is the domain of the experiment. On the theoretical side, the waveform has to be predicted in order to be recognized as a signal, as opposed to noise. This prediction is possible by solving the equations of Einstein's general theory of relativity, which is not a trivial task. It is done by computer modeling, with the help of methods such as post-Newtonian analytic approximation and numerical relativity. The data from the interferometers is compared with a bank of theoretically predicted waveform templates in order to find the one that best matches the observed data, which is a method called matched filtering (Abbott et al. 2016a). Without the experiment, no observation is possible, while without the theory, it would not be known what there is to be observed. Since it was possible to model the signal from binary mergers because of the equations, these signals were detected first.

“According to general relativity, a pair of black holes orbiting around each other lose energy through the emission of gravitational waves, causing them to gradually approach each other over billions of years, and then much more quickly in the final minutes. During the final fraction of a second, the two black holes collide into each other at nearly one-half the speed of light and form a single more massive black hole, converting a portion of the combined black holes' mass to energy, according to Einstein's formula $E=mc^2$. This energy is emitted as a final strong burst of gravitational waves. It is these gravitational waves that LIGO has observed.” (LIGO 2016, 1-2)

Of the three possible binary coalescences, it was not clear which one would be detected first.¹⁶ Some sources published before the detection point to the black hole mergers as the least

¹⁶ Let me pause here and reflect on the jump from talking about a detection of a gravitational wave signal to talking about the detection of a gravitational wave itself, and even further, to talking about a detection of both the signal's and the wave's source. It is often said that the detections allow us to “hear” or “look deeper”, or “open a window” to the universe. The waves are considered to be directly detected and hence observed by a detection of a particular signal by an interferometer that matches a predicted signal. This signal represents the wave with its distinct features, whose details make it possible to learn about the source of the wave, in this case binary collisions, most often black hole mergers. Since these events happened billions of years ago, this is the only way we can “see” them. Therefore the mergers themselves are also considered to be detected and observed. Although it seems counterintuitive to directly observe something that has happened in the past, when what is observed is the information about a past event mediated by its effect, for the sake of simplicity, convention, and continuity we will accept

understood of the compact binaries, therefore least likely to be detected first. In Harry Collins' 2004 book *Gravity's Shadow* we can still read that "inspiraling binary black holes should emit a strong signal but have a waveform that as yet cannot be fully calculated" (p. 662). Apparently there was no consensus on which kind of binary events were to be detected first, which I was also told by a researcher in gravitational wave physics.¹⁷ In her report black hole mergers, however, fared much better:

"In the beginning we predicted that we will... well, there are some people who believed that neutron binary star will be detected first and other people believed that binary black hole will be detected first." (Scientist I)

Not only that the signal from a black hole merger was detected first, but they have also been the most frequent sources of detected gravitational waves so far. Together with supernovae, binary mergers are the loudest events in the universe. In fact, supernovae are considered to be the loudest. How is it then that a supernova has not been detected first? It has to do with the impossibility of modeling the signal from a supernova, since too many uncertainties are involved. Black holes and neutron stars are simply more calculable, which makes them the "easiest" targets, and most proximately achievable goals. An additional problem is that supernova events are expected to occur rarely, at least those close enough to be "heard" by the

that the sources of gravitational waves have also been detected and observed. This is the terminology often used by researchers themselves, and this thesis does not aim to provide a discussion of observables and unobservables, least their relation to the ontology of scientific theories.

¹⁷ I am grateful to three gravitational wave scientists who agreed to answer my questions about the gravitational wave physics and their own research. The researchers are members of the Max Planck Institute for Gravitational Physics in Hannover called the Albert Einstein Institute (AEI). The AEI closely collaborates with the LIGO and is a part of the LSC. It comprises two Institutes: AEI Potsdam and AEI Hannover. AEI Hannover operates the GEO600 detector as well as the most powerful computer cluster dedicated to gravitational wave data analysis, the Atlas. One of the interviewed scientists is a research group leader, one is a postdoc, and one is a PhD student. When needed, they will be called Scientist I, II, and III, a notation that follows the occurrence of their reports in the text and not necessarily the aforementioned sequence with the specification of their academic status, in order to preserve their anonymity. This notation will only be uttered when there are no specifications of their narrower field of work or academic status.

current detectors – an estimated 1 per 30 to 50 years.¹⁸ With supernovas, however, the prospects of unexpected findings that would expand our knowledge are the biggest, because it is not known what to expect, hence “searching for burst gravitational waves requires being utterly open-minded” (LIGO: “Sources and Types of Gravitational Waves”, n.d.). However, if a supernova would explode symmetrically, no gravitational waves would occur, since they can be emitted only by asymmetric accelerating bodies.

There are two other types of gravitational waves: continuous and stochastic. Continuous waves are emitted from single spinning bodies, most likely from neutron stars. Again, perfect spherical shapes will not emit gravitational waves, but in the case of a neutron star, “the bumps and imperfections will generate continuous gravitational waves as it spins” (LIGO: “Sources and Types of Gravitational Waves”, n.d.). Since the spin rate is constant, so are the gravitational waves emitted constantly, with the same amplitude and frequency, and lasting for a long a time (ibid.). However, they are very weak, “quiet”, so it is harder to detect them, and they have to be integrated over a long period of time, as opposed to transient events like binary mergers. Continuous gravitational waves have not been directly detected yet, although they can be modeled. Stochastic gravitational waves are the smallest, the “quietest”, and therefore the most difficult waves to detect. They are mixed gravitational signals coming from every direction and distance and interfering with one another randomly (stochastic), which can be analyzed only statistically, but not modeled precisely (ibid.). These signals might even include relicts from the Big Bang, but they are considered to be far below the sensitivity of the existing detectors (Meadors 2014, 9-10).

The scientific objectives of the LIGO include “research in the fundamental physics of gravitation as well as in astronomy and astrophysics” (LIGO 2018, 4). Fundamental research is understood here in terms of goals, such as:

“(…) further detections of gravitational waves, test of General Relativity in the strong field and high velocity limit, direct measurement of the polarization and propagation speed of gravitational waves, direct observation of the dynamics of black holes, and constraining the neutron-star equation of state” (LIGO 2018, 4).

¹⁸ As reported by Chris Pankow, a postdoctoral fellow at Northwestern University, Center for Interdisciplinary Exploration and Research in Astrophysics, in a discussion in *LIGO Magazine*, March 2019, p. 8.

Besides stated scientific objectives, technical objectives include upgrading the sensitivity of the current detectors and building the Indian detector, which would provide a triad of detectors separated by intercontinental baselines. Broader societal impacts include education, public outreach, and technology development (LIGO 2018, 4-5). There are multiple technological advances resulting from LIGO's and LSC's activities: the method used to stabilize LIGO's laser frequencies is also used to build semiconductors in computers and cell phones, while other spin-offs are being applied in measurement science, seismic isolation, vacuum technology, mirror coatings and optics (NSF 2017).

LIGO is a unique project in terms of ambition, risk and long-term orientation. Richard Isaacson, a retired NSF Program Director for Gravitational Physics, wrote about it:

“LIGO was a project initiated long ago, at a time when the scientific community defined basic research priorities at NSF in a bottom-up fashion. Then, the function of NSF was to help scientists do what they found interesting. NSF believed that basic research belonged at universities, as you could never predict what the outcome would be here. LIGO was the result of investment in long-term research, development, and construction (over four decades – a scientific lifetime) to reach towards a difficult but enormously exciting goal. During that period, the NSF oversight philosophy was to get good people and try to stay out of their way as much as possible, but to stand by and be ready to help with mid-course corrections when needed. During much of this time, Congress was ready to take risks to achieve significant progress, and to show patience when things hit a bump.

LIGO was conceived at a time when scientists, administrators, and politicians showed great vision. We now live in a different era. There is confusion about the value of basic research compared to applied science and engineering. Key Congressional committees are led by politicians who do not believe in evolution.¹⁹ Short term goals are important, and long term vision is rare. A similar project with such high levels of risk could not be attempted today. It is important that LIGO achieves major successes and solves many

¹⁹ I am inclined to think that what is meant here is that key politicians do not believe in the evolution of scientific projects. They do not believe that it is better to let research develop in unexpected ways over a longer period of time, rather than to target it to specific short-term goals. Another possible reading suggests that key Congressional politicians do not believe in Darwin's evolution of species, but I do not have enough information to back it up, and it is irrelevant for this thesis.

cosmic mysteries, to repay the trust and commitment of a public that has invested much and waited patiently for a long time to see it operational.” (Isaacson 2016, 61)

The view of basic science advocated here accords with the view put forward in Vannevar Bush’s 1945 *Science, the Endless Frontier*, often called the linear model of science.²⁰ In the linear model, unconstrained scientific freedom is believed to be the best way to provide epistemic insight which will then be followed by useful applications. Hence the linearity – first comes the basic science, primarily in the academic context, then the applied, mostly in industry.²¹ Furthermore, basic science is understood as open-ended in terms of both the timeline and research goals. Therefore the “long-term vision” as opposed to short-term goals. It would be hard to convince anybody to fund LIGO had the funding been conditioned upon short-term goals.²² The linear model has dominated US postwar science policy until the last few decades, when the need to apply, and more recently, to translate, became much more pronounced, and the traditional domains of academy and industry have been broadened and started overlapping. However, although LIGO has been initiated in the late 1970s, by 1992 when the funding for the initial construction was approved, the shortcomings of the linear model were well recognized and it had already ceased to be the dominant view, with emergentism and interactionism taking over (see Adam, Carrier, and Wilholt 2006; Wilholt 2006; Carrier and Finzer 2011). The arguments against the linear model emphasize that innovation is in fact complex, uncertain,

²⁰ Philip Kitcher calls it the “seed-corn argument” (2011, 123).

²¹ The linear model is sometimes distinguished from the cascade model of science. According to the cascade model, the applied is logically dependent on the basic. In the linear model, the transformation follows in a temporal sequence, not necessarily a logical one. The idea is that “it is of no use to attack a practical problem by research narrowly targeted at this problem” (Carrier and Finzer 2011, 86). The cascade model allows for fundamental insights occurring in the course of practical projects, so it does not need the temporal priority of basic research. In Adam, Carrier, and Wilholt (2006) the cascade and the temporal model are discussed as the same model (p. 437). However, both of these models are outdated today and different forms of interactionism and emergentism have taken over. More on these models will be said further in the text. For an overview and discussion of different models of science, see Carrier and Finzer (2011) and Adam, Carrier, and Wilholt (2006).

²² However, the technological development that has enabled the detection is certainly a domain of application, though it has never been the primary research goal. In fact, Collins argues that the expected spin-offs in laser technology and seismic isolation were an important argument in favor of funding LIGO (2004, 489-511), and they were definitely achieved on a shorter timescale than the detection itself.

difficult to measure, dependent on the market environment and social context, and occurring in both basic and applied research (see Kline and Rosenberg 1986 and Sarewitz 1996, 97-117). Interactionism highlights feedback loops between basic and applied science, while emergentism sees applied research and technology development as independent from basic science, providing its own solutions to particular, local problems. Interactionism accounts for the cases in which general rather than local understanding is produced in the course of solving practical problems, because fundamental questions can arise in very specific local research. Interactionism still acknowledges the practical usefulness of theoretical knowledge which is often adapted for the needs of local problems, in accordance with the linear model. Since it allows for genuine insight occurring in applied contexts, it is a middle position between the linear model and emergentism. In Adam, Carrier, and Wilholt (2006) the interactionist position that is put forward is called moderate emergentism, while the linear model is absorbed by the cascade model:

“On the one hand, it is granted to emergentism that applied research may give rise to genuinely new knowledge that is adjusted to the relevant local circumstances. On the other hand, it is conceded to the cascade model that the knowledge gained is not purely local. Rather, it transcends the specific conditions that make up the practical challenge at hand. (...) Technological innovations tend to rely on comprehensive theoretical principles and may entail new theoretical insights.” (2006, 443)

LIGO was initiated at a time when academy was already stepping in the era of technology transfer,²³ and it has been funded for construction in the wake of translational research initiatives, to which I will dedicate the second part of the thesis. This suggests that there had been a rationale for funding LIGO beyond the mere linear model instantiation, although LIGO is a clear case of basic science conceived in a “bottom-up fashion” where the direction is from the scientists to the policy makers and to the broader society. Researchers led by scientific curiosity set the priorities, while policy makers “stay out of the way” “show patience”, and are “ready to help”. In contrast, the recent model of “Responsible Research and Innovation” with the leitmotif “Science with and for the society” is characterized by the exactly opposite direction

²³ The US Bayh-Dole Act of 1980 permitted universities, small businesses, and non-profit institutions to own patents, as opposed to assigning them to funders, usually the government. This change allowed universities to enter the market economy and to transfer new technologies to real world contexts (see Loewenberg 2009).

of setting research priorities than the one exemplified in LIGO. In the new approach, science is informed by the needs of the society and is enriched by the plurality of diverse non-expert perspectives. It can support long-term research projects as long as they are responding to an ongoing need like fighting cancer or dealing with climate change. LIGO is not such a project in terms of direction, since there is nothing socially relevant to respond to by setting it up, but the long-term orientation can still be a shared feature. What LIGO provides is new understanding, and the risk is mainly posed by the prospect of investing so much at the expense of other research directions, and possibly not delivering on the promises. However, basic science still gets funded both publicly and privately (see Grant 2017), and it is hard to speculate would it be the case that a similar project could not be attempted today, or how we would be assessing LIGO had it not have resulted in success.

But LIGO has indeed become operational. Ten binary black hole mergers and one neutron star merger have been observed in the second observation run (O2), while five preliminary reports of the detections have been released in April 2019, following the start of the third observation run. As the detections in O2 were happening, there was still pessimism even on the part of the members of the gravitational wave community with regards to the timeline of the research and the achievement of research objectives:

“Three hours before the first binary neutron star merger was detected, I was at a conference listening to people say we might never see enough events to provide useful insight. When I got off the plane in Berlin several hours later, I saw an email about the first detection in conjunction with a gamma ray burst, and we soon knew that we might eventually see hundreds. Science often feels like it moves slowly, but occasionally an entire field can change in a day.” (Benjamin Lackey, Senior Scientist, Max Planck Institute for Gravitational Physics)²⁴

The first detection of a gravitational wave from a binary neutron star merger happened on August 17, 2017, and the signal was followed by an electromagnetic counterpart in the form of a gamma ray burst. One of the scientists I talked to emphasized a particular role that this detection had:

“The first neutron star detection was important because that gives us a link to the rest of astronomy and makes us, you know, really like astronomers.” (Scientist II)

²⁴ Max Planck Institute for Gravitational Physics (AEI), 2017 (“Researchers’ voices”)

This is so because astronomy has up to now been mostly focused on detecting electromagnetic radiation, so the occurrence of an electromagnetic signal together with a gravitational wave signal links the two research fields tightly together.

While we are still talking about gravitational wave physics, what we are witnessing in real time is its transformation into gravitational wave astronomy, as more facts about the universe come to be known through “messages” carried by gravitational wave signals. Following this transformation, the astronomy itself is transforming into multi-messenger astronomy and multi-messenger astrophysics, as knowledge is being acquired through a variety of signals detected by different telescopes and observatories, and then measured and interpreted in coordination and combination. I will thus start the analysis of time constraints in science with this fascinating research field where “science often feels like it moves slowly”, in the heart of fundamental research – in gravitational wave physics.

2. Speed of research and goal setting

It has been said in the introduction to this thesis that time constraints can mean different things: external events that need to be urgently addressed, deadlines posed by funding agencies, self-estimated milestones, particular time slots in which an event of interest occurs and can be observed or analyzed, or an individual lifetime in which a certain epistemic goal is pursued. More generally, time is often described under the heading of pragmatic factors, together with other resources like money or computational power. It is definitely limited, no matter what the reference frame is. We will now focus more closely on LIGO and LSC, given the unusually long-term orientation that their research has had from its inception, in order to spell out the ways in which time has nonetheless been optimized by setting particular epistemic goals and by designing methods to achieve these goals. It would be hard to imagine that scientists have just been sitting around and waiting for an extraordinarily loud event to surprise them. They have indeed been waiting for a particular event, but they have also been constantly improving their chances of coming across one.

Let us first take a look at the LIGO timeline in order to grasp the milestones on the historical scale. This is the timeline of the LIGO research as presented on the LIGO webpage, starting from the pioneer work on gravitational wave detection by means of laser interferometry in the 1970s, until the 2017 Nobel Prize for the first detection (LIGO: “Timeline”, n.d.):

1970s – Early work on gravitational-wave detection by laser interferometers, including a 1972 MIT study describing a kilometer-scale interferometer and estimates of its noise sources

1979 – National Science Foundation (NSF) funds Caltech and MIT for laser interferometer research and development

1983 – MIT and Caltech jointly present results of the kilometer-scale interferometer study to NSF. Receive NSF committee endorsement on new large programs in physics.

1984 – LIGO founded as a Caltech/MIT project. National Science Board approves LIGO development plan.

1986 – Physics Decadal Survey and special NSF panel on gravitational wave interferometers endorse LIGO

1990 – National Science Board (NSB) approves LIGO construction proposal, which envisions initial interferometers followed by advanced interferometers

1992 – NSF selects LIGO sites in Hanford, Washington, and Livingston, Louisiana. LIGO Cooperative Agreement signed by NSF and Caltech.

1994-95 – Site construction begins at Hanford and Livingston locations

1997 – The LIGO Scientific Collaboration (LSC) is established and expands LIGO beyond Caltech and MIT, including the British/German GWO Collaboration, which operates the GEO600 interferometer in Hannover, Germany.

1999 – LIGO inauguration ceremony

2002 – First coincident operation of initial LIGO interferometers and the GEO600 interferometer

2004 – NSB approves Advanced LIGO

2006 – LIGO design sensitivity achieved. First gravitational wave search at design sensitivity. Science Education Center inaugurated at the LIGO Livingston Observatory.

2007 – Joint data analysis agreement ratified between LIGO and the Virgo Collaboration, which operates the Virgo interferometer in Cascina, Italy. Joint observations with enhanced initial LIGO interferometer and Virgo.

2008 – Construction of Advanced LIGO components begins

2010 – Initial LIGO operations conclude; Advanced LIGO installation begins at the observatories.

2011-2014 – Advanced LIGO installation and testing

2014 – Advanced LIGO installation complete

2014-2015 – Advanced LIGO sensitivity surpasses Initial LIGO

Sept 14, 2015 – Advanced LIGO detects gravitational waves from collision of two black holes

Dec 26, 2015 – LIGO observes a second binary black hole coalescence

Jan 12, 2016 – First observing run of LIGO advanced detectors ends

Jan 4, 2017 – LIGO observes its third binary black hole coalescence

Aug 14, 2017 – Gravitational waves from a binary black hole merger observed by LIGO and Virgo

Aug 17, 2017 – LIGO and Virgo make first detection of gravitational waves produced by colliding neutron stars

Aug 25, 2017 – Second observing run of LIGO advanced detectors ends

Oct 3, 2017 – LIGO co-founders Rainer Weiss, Barry Barish, and Kip Thorne are awarded the 2017 Nobel Prize in Physics

We can see here the pace of progress in gravitational wave research via laser interferometry, but it is important to note that laser interferometry had not been the first or the only attempt to detect gravitational waves. Earlier attempts include room temperature resonant bars (early 1960s to 1975), resonant spheres (1970s -1990s), and especially cryogenic bars (mid-70s to now)

(Collins 2004, 2017).²⁵ Laser interferometry and cryogenics have been the only means of gravitational wave research that has eventually survived in terms of funding, based on the evaluated prospects of reaching the levels of sensitivity needed for an actual detection. Though cryogenic bars have been mostly abandoned as a means of detection, cryogenically cooled mirrors are being used in the KAGRA detector in Japan.²⁶

Taken all the attempts together, gravitational wave research has taken more than fifty years until the first detection, and it continues to develop a new era of multi-messenger astronomy. Is 50 years long? What we are interested in is evaluating whether the cognitive, material, temporal, and other resources have been optimized in order to bring about the epistemic advancement that has been sought for. Could gravitational waves have been detected earlier? Which initial goals had been specified and why? Did considerations about the time of the detection bear on decision making about goal setting and research methodology? Which methods have been developed to reach the goals?

The benefit of hindsight gives us reasons to think that resources have indeed been optimized to a sufficient degree. What we can point to are certain methodological choices that have been made in order to achieve the epistemic goal as soon as possible. Moreover, the choice of a particular epistemic goal also reflects considerations about the time of the detection. Of course, delays, errors, and mishaps permeate any practice, scientific practice not excluded, but the research on gravitational waves has undoubtedly been improving and making progress during the years, even without actual detections taking place. The progress was made through upgrades in detector sensitivity, waveform modeling, and by setting upper limits on the flux of gravitational radiation.²⁷ It is hard to estimate exactly how optimal these advances have been

²⁵ Harry Collins' *Gravity's Shadow* (2004) provides a detailed insight into the history and sociology of gravitational wave physics, with an especially thorough account of earlier attempts to detect gravitational waves, and the reasons for the prevalence of laser interferometry.

²⁶ In 2015, two cryogenic bars were still operating in Italy (Collins 2017, 52). LSC Program for 2018-2019 mentions cryogenic interferometers as "an attractive approach to lower the test mass thermal noise" (LSC 2018, 18), while David Shoemaker from the MIT, the LSC spokesperson, said that "the use of cryogenics, in particular, might be essential if future detectors are to offer vastly improved sensitivity" (Castelvecchi 2019, 10).

²⁷ An upper limit states how many events of a particular kind can be expected. According to Collins (2004), it can state the following: "The flux of gravitational wave is less than n bursts, of strength y , within frequency band x , per day/week/year", or "The strength of continuous gravitational waves within

or should have been, but the project has successfully survived the challenges and exemplifies an achievement of a goal set on a large time scale.

There has always been a recognition of the element of luck that will accompany the detection if it is to be considered real and announced to the wider community. This is so because the sensitivity of ground detectors at any particular time will be good enough for detecting events of a certain frequency band, at a particular distance, occurring in a particular portion of the sky. Although we are “bathed” in gravitational waves every day, only the ones “loud” enough or whose sources are close enough will have the chance to be detected with the required level of significance. Signal to noise ratio is a measure of whether a particular signal has been detected at a statistically significant level, where low values mean that a signal has not been detected, while very high values imply a nearly certain detection (LSC and Virgo 2017, 2). The first signal, GW150914, has been observed with a signal-to-noise ratio of 24 and a false alarm rate estimated to be less than 1 event per 203.000 years, which is a significance greater than 5.1σ (sigma) (Abbott et al. 2016a, 1). As a result of gradual advances over the years, and especially the improved sensitivity of the Advanced LIGO, the detection has been expected with a higher probability as compared to the first observing run. Following the experimental and theoretical upgrades, more detections have been expected after the first one, and indeed, the Advanced LIGO-Virgo detection rate for black hole binaries has up to now been 1 detection in 15 days of observing time on average.²⁸

Before going further into goal setting and methods used, I will first argue against characterizing any particular research practice as “slow” or “fast”. At least two reasons can be offered. First, it is often unclear how to establish definite starting and ending points of a particular research, since every instance of research depends on the research that preceded it or contributed to it. In the gravitational waves case, laser interferometry existed before the LIGO collaboration was established. LIGO interferometers are essentially Michelson interferometers, a device invented in the 1880's by Albert Abraham Michelson, before the theory of general relativity was formulated (LIGO: “What is an Interferometer?”, n. d.). They are, of course, adapted for the search for gravitational waves, but it is nonetheless hard to estimate the time it would have

frequency band z is less than p .” (p. 699). See Collins (2004, 698-726) for a discussion of setting upper limits as a valid scientific result in the absence of actual detections.

²⁸ As reported by Thomas Dent, a group leader in the Galician Institute for High Energy Physics in Santiago de Compostela, in a discussion in *LIGO Magazine*, March 2019, p. 7.

taken for the detection to happen had it first been necessary to invent interferometry from scratch. But hardly anything is science starts from scratch, since the material and cognitive tradition is quite substantive.

Another reason is that the fact that a research project takes more than fifty years can mean that it is slower or faster than another project, but only in comparison, since slow and fast are relational terms. There is no independent way to assess the standard of the research speed in achieving goals, and we have to stay ignorant as to whether the detection happened slow or fast. Taken the complexity in question, maybe it came extremely quickly. We will soon see how the decision to build two interferometers as opposed to one was what made the detection possible in the first place, and what also made it relatively temporally proximate. Moreover, the decision was made in the face of great opposition, especially because it required more resources.

In the conclusion of his 2004 book, Harry Collins refers to the LIGO timeline as short:

“Again, the stresses of growth from small to big science are made visible in a peculiarly dramatic way, because LIGO is making this change for the first time and over a short timescale.” (p. 790)

Fairly enough, Collins refers to the sociological transition of gravitational wave science in terms of organizational and management style. The epistemological interest in the LIGO timeline that motivates the analysis in this dissertation is better captured on another place in his book:

“As things are, gravitational wave detection, from instigation to first detection, is a rough match for an adult lifespan or professional career (including that of this author), and it would be hard to persuade scientists to devote their lives to a scientific enterprise with a significantly longer time span. As Rich Isaacson of the NSF put it to me [1995] “Once you... can run in coincidence, you can do science. People are not going to spend ten years of their lives just building technology. I think that’s the real reason we don’t just build one.” (Collins 2004, 773)

Fifty years of research is long from the perspective of a single researcher who might not live to see the results. But what is it that makes it acceptable and even promising? What is problematized here is the aforementioned decision to build two interferometers as opposed to one. One way to deal with the noise in the detector is to discount it by cross correlating two devices, as it was mentioned earlier. Each detector will pick up a great deal of noise but only two will simultaneously record a gravitational wave signal. This is a very elegant way to look

for signals from burst and binary inspiral sources. It is, however, not as helpful when it comes to continuous waves, since they might in principle be detected with only one detector if the frequency of the waves is known, because they run continuously but are very quiet, hence buried in the noise (Collins 2004, 662; Pitkin and Sun 2017, 10-12).

Both of these decisions: to focus on the burst and inspiral sources as the first detection target, and building two interferometers to facilitate the search for these sources particularly, were decisions where pragmatic considerations about the time of the detection certainly played a role. But the fact that they can run in coincidence opens up the very possibility of a detection, no matter how initially small, and this is epistemically promising, as opposed to “just building technology”, which is practically impossible to be epistemically fruitful in the same time span. It means that the epistemic underpinnings of the decision to build two detectors have been crucial, but also inseparable from pragmatic ones. The decision was not based only on concerns about the researchers’ timelines and maintaining funding, it was based on the prospects of a successful detection. In other words, it was about achieving epistemic goals, and about achieving them sooner rather than later. Although it takes more time and money to build two detectors as opposed to one, the trade-off is acceptable taken that two would make an easier, faster, and more convincing case for distinguishing the actual signal from noise, which means that both time and money are likely to be saved, but most importantly – that the epistemic goal has realistic prospects of being achieved.

Scientists firmly argue that pragmatic considerations involved in building two detectors were not employed for the sake of economizing resources of time and money, but that they were employed to bring about an epistemically superior method:

“I argue in *Gravity’s Shadow* that while one detector would be enough to do scientific development work, it was much wiser for various human and economic reasons to build two, but the scientists I argue with invariably tell me: ‘we had to build two or we would not have had coincidences and would not have been able to see anything.’” (Collins 2017, 55, footnote 2)

We find here the sociological argument – scientific justification is all about human and economic reasons, contrasted with the epistemological argument – scientific justification is all about epistemic reasons. My analysis pertains to a third conclusion: the two are inseparable, especially when it comes to the relation between epistemic goals and time constraints. We do not have unlimited time to work on an idea. We adjust the goals and methods to the time we

can reasonably expect to have. In this particular case, it eventually meant building two interferometers in order to search for gravitational wave signals from inspiral and burst sources. The resonant and cryogenic bars failed; one interferometer, even when eventually sensitive enough, would not be able to single out a particular signal as a candidate without extensive analysis, which might stall the research and rarely, if ever, achieve a sufficient significance level for the announcement of a detection. A second interferometer allows for a double check or a version of a replication: if there is a gravitational wave, it will be detected on both sites.

Of course, the research was not going as smooth as that. LIGO has broken a few estimated deadlines and some bets with regards to the time of the detection had been lost (Collins 2017, 56), but nevertheless, choices that had to be made were made for good reasons. Most importantly, these reasons have been both epistemic and pragmatic, i.e. truth attaining, but temporally constrained, primarily by funding cycles. They had to make realistic predictions of the time of the first detection in order to convince the community that it is something worth doing, both scientifically and financially. Apart from the decision about the experimental method that should be employed, the theorists have constantly been modeling ever new predictions of the waveform, in order to be ready when LIGO goes online and starts commissioning data. In fact, it was the theoretical predictions that granted success to the project. The predictions of Einstein's general theory of relativity have been repeatedly confirmed by observational and experimental methods, such as the perihelion advance of Mercury,²⁹ the time dilation of global-positioning satellites,³⁰ and the bending of light around big astrophysical objects (Ohme 2012, 1).³¹ The predictions concerning the existence of gravitational waves, black holes, and some of their features were therefore also undoubted, but yet observationally unconfirmed. As Collins puts it:

“(...) the willingness of the scientists and the funding agencies to press ahead on the basis of the theoretical certainty that there really was something out there to be seen, and

²⁹ Mercury's perihelion, i.e. the point on its orbit when it is closest to the Sun, does not follow the same path each time, but it departs or advances from it over time due to the curvature of spacetime.

³⁰ GPS satellites are in orbits high above the Earth, where spacetime is less curved than on the Earth's surface, so their clocks seem to be ticking faster than those on the ground when viewed from the Earth. This is because time appears to be moving slower close to massive objects. The GPS satellite clocks thus have to be set to account for the difference.

³¹ This happens due to the deflection of light rays close to a spherical mass which curves the spacetime.

that the next generation of detectors should see it, is a great triumph of human perseverance.” (Collins 2017, 56)

The Initial LIGO observing run lasted from 2002 to 2010 and no detections were made. After four years of engineering improvements, the new Advanced LIGO started its first observation run in September 2015 and within days made its first detection of gravitational waves. The latest observing run that started in April 2019 has resulted in five preliminary reports of the detections in April only. What can we say about this particular timeline: that the first eight years was long? That “within days” is fast?

When asked about the timeline of the gravitational wave research, i.e. prior to any detections, what were his thoughts about the possible time of the first detection, one of the gravitational wave researchers in Hannover said that he believed that the detection would occur later, that it would take longer than it actually did. He also shared that opinions on the time of the detection were different among members of the research community, and that some were convinced that it was “around the corner”, while others were more skeptical. He also said that he did not really care when the detection will happen:

“It’s a fascinating subject to research and, you know, if Nature is kind and bla bla bla we will be able to do it soon, but if not, the detection will come later. From a purely idealistic point of view, I didn’t care too much.” (Scientist III)

Nature is often invoked in the reports of gravitational wave physicists:

“The combined measurement of a gravitational wave signal, a gamma ray burst, and a kilonovae is a generous gift from Nature, I am glad that my pessimistic expectations have been proven wrong.” (Tim Dietrich, Postdoctoral Researcher, Max Planck Institute for Gravitational Physics)³²

“Nature had indeed surprised us! In a matter of hours, analysts started analyzing the event, the detector characterization group began looking at data artifacts that could rule it out, instrumentalists wondered who could plant a double blind injection in the detectors, and the LSC management scratched their heads to figure out the next step: there was no room for mistakes, we had to be sure this was a detection before we could announce it to the world.” (González and Cavaglià 2016, 5)

³² Max Planck Institute for Gravitational Physics (AEI), 2017 (“Researchers’ voices”)

“It does justice to all the past years of hard work and non-detections and to this wonderful gift of Nature that GW150914 has been.” (Maria Alessandra Papa, Senior Research Scientist, Max Planck Institute for Gravitational Physics)³³

Nature’s “kindness”, “surprise”, or the “gift” that has been given is the occurrence of a gravitational wave signal so “loud”, “clear”, and “near” that the current instruments were able to detect it and the current theory able to understand it. Simultaneity of these factors make for a successful discovery: 1. a particular event in the empirical world – a gravitational wave produced billions of years ago passes the Earth at a point in time; 2. a well calibrated instrument at the appropriate sensitivity level measures it; 3. a theory explains it – Einstein’s general theory of relativity first presented in 1915 and published in 1916; and 4. computational and analytical resources are able to cope with the uncertainties involved in modeling the signal from the equations provided by the theory.

It is important to emphasize again the role that simultaneity, synchronicity, or “running in coincidence” plays in this particular case. Timely advancement in this fascinating field is to a large degree conditioned by successful synchronization of different aspects of theory and reality. The synchronization of theoretical and experimental research is a structured effort to allow for a serendipitous discovery – serendipitous because of the coincidence with a particularly loud event, such that allows for “surprises” by both its occurrence and unexpected properties, but otherwise carefully attuned to a particular epistemic goal.³⁴ The broader take home message is that the pragmatic and the epistemic aspects of research are inseparable. Burst and inspiral sources are the first targets because there has been an epistemically promising way to go about them – building two interferometers that run simultaneously. Does it take longer to build two, or to improve one until reaching the necessary sensitivity? Deciding for the synchronicity of the two interferometers was an epistemically superior option, and we have good reasons to think that it has also saved time and enabled further development of the field.

³³ In “A Perfect Source: Timeline of GW150914”, *LIGO Magazine*, March 2016, p. 12.

³⁴ I adopt an understanding of serendipity as advanced by Samantha Copeland (2015, 2017) according to which a serendipitous discovery includes chance, sagacity, and a valued outcome, and is “both unpredictable and yet can be cultivated” (Copeland 2017, 1). I will say more about her notion of serendipity in Part II, especially in section 6.1., p. 112-113.

3. Prioritization and urgency

The scientist's point of view ("From a purely idealistic point of view, I didn't care too much.") is very much opposed to the perspective of the NSF program director Richard Isaacson, who finds it important that LIGO "repays the trust and commitment of a public that has invested much and waited patiently." The clash between the scientist's and the policy maker's perspective is not particularly surprising, but what is philosophically interesting are possible pragmatic considerations in the process of knowledge production, such as considerations about the time of the occurrence of results, and how these are reflected in the prioritization of research goals and in the design of research methods. We can find such pragmatic considerations in the following report of a researcher working in the field of continuous gravitational waves, which have not been detected yet:

"We'd trusted the calibration from the beginning, but it's a little bit different when you really start seeing things for the first time. So that does make us more confident that what we're doing is worth, is possible, and is worthwhile, but we're still not sure when we will make our first detection... It's theoretically imaginable that we could see something next year, but it could not be for another decade.... And that's, that makes me think of my own time in my own carrier process."

Continuous waves are held to be a promising source because they are present in the LIGO data all the time so one does not need to be lucky to observe them at exactly the right time, but the problem is that the data processing techniques are still inadequately sophisticated to discern the signal from noise. Although it is a case of epistemically promising research, pragmatic considerations with regards to speed of getting results enter the scene. It may be more fruitful to move somewhere where foreseeable results are more likely to obtain. For the continuous wave researcher, the temporal proximity of the occurrence of future detections is a source of worry when it comes to the choice of further projects to work on inside the same research community. It makes sense to ask oneself how to proceed with research or which research to proceed with, taking into account when the results will obtain and what could have been done in the meantime if they had not obtained, i.e. what is lost in terms of opportunities. The worry of our continuous wave researcher makes sense because we should aim at maximizing gain and minimizing loss, both epistemically and non-epistemically, and this requires trade-offs and prioritizing.

In more general terms, prioritization involves decision making at the intersection of epistemic, pragmatic, and ethical considerations. Apart from epistemic gains, there is no pressing need on the societal level for the detection of gravitational waves, or at least we cannot conceive of one as pressing as that of Ebola virus at one time, Zika virus at another, or that of continuous sources of worry like cancer, cardiovascular diseases, depression, and a cluster of humanitarian, health and environmental issues stemming from conjoined problems of social inequalities and climate change. To these problems we attend not because the results are foreseeable, but because they are strongly needed. In other cases, the results are wanted as a return on investment – of funds, people, education, infrastructure, time, and experiments, as it is in the gravitational wave research. There is a loss of resources when results do not obtain.

Focusing on transient events, primarily the binary inspirals that can be modeled, the epistemic goals of the gravitational wave community have been set: they want to learn as much as possible from these sources and they are gradually getting exactly where they aimed to be, since the number of detections is increasing with the increased sensitivity of the instruments. Not only would it be inopportune to allocate significant resources to searches for continuous waves, it might also be compromising the pace of progress in the field. New detections bring new knowledge and set new priorities. If the algorithmic searches for continuous waves surprise us, all the better, but that is now, if it already has not been, a maverick way to go about gravitational waves.³⁵

“A central question of the waveform modeling community to prepare for the era of advanced detectors is how the limited computer resources should be spent most efficiently? (...) we point out that the model errors we find are potentially good enough for the first detections and interpretation of signals whose amplitude is close to the detection threshold. In that sense, the prospects of being prepared for immediate astrophysical applications of GW detections are rather good, provided that our models are subsequently refined (...) A reliable method to address more fundamental questions,

³⁵ I will say more about the methods used in searches for continuous gravitational waves in the next chapter. See Kitcher (2011, 206-208) for a discussion of the division of cognitive labor, especially conformist vs. maverick strategy. Conformists adopt a strategy of working on the most promising research in a particular field, where consensus is established and advances are often small but unchallenged. Mavericks, on the other hand, work in a not so well established research area in the field, where success is not guaranteed or even likely, but in case it does occur, the reward is much higher.

such as testing General Relativity or determining the nature of matter in neutron stars from GW observation, however, requires much more accurate models than those we can construct today. (Ohme 2012, 3)

Note that this was written before the first detection, but it makes clear that the more detections there are, the more of subsequent refining is possible, hence more accurate waveforms can be modeled in the future. Once the first detection is achieved, the information gained by it shapes and informs the direction of further research. In this case, insisting on less probable or fruitful research directions, though possibly resulting in success, does not optimally capitalize on the new opportunities set by the actual detections, where it makes sense to expect a gradual progress.

“Very loud events are expected to occur extremely rarely with the current sensitivities of ground-based detectors, and the success of these instruments crucially depends on the ability to detect as many signals as possible, particularly close to the detection threshold (...) therefore, good waveform models that resemble the true GW signals with a very high match are invaluable.” (Ohme 2012, 27)

Clearly, what has been prioritized are the inspiral binary searches, and more recently, black hole binaries specifically. First, because they can be cross correlated by matched filtering, second, because they can be modeled, and third, because they have been repeatedly detected, hence setting the path for further advancement. Each event changes the goals and redistributes resources. Black hole binaries have already opened up a new set of possibilities simply by allowing a catalogue of detections by now. In this new context, new goals are being set.

“Perhaps the most important open question for black hole mergers is: what’s the distribution of the actual properties of the sources? This is something that we’re only beginning to scratch the surface on – just recently we published our catalog of detections and performed the first population predictions for binary black holes with ten binaries! Of course, ten events may seem like a lot now, but it’s insufficient to pin down more than a few major points.” (Chris Pankow, a postdoctoral fellow at Northwestern University, Center for Interdisciplinary Exploration and Research in Astrophysics)³⁶

Ten detections is a lot when the goal is any detection. But ten is not enough if something is to be learned about the population of these sources in the universe, which is a goal that is being

³⁶ From a discussion in *LIGO Magazine*, March 2019, p. 8.

developed, only “beginning to scratch the surface on”. New goals similar to that, contingent on what the next discovery or a cluster of discoveries will be, are being set now, because after the detection the field changes entirely. It becomes less clear on what to focus on and which research directions to prioritize.

“Previously it was very, very focused and one could easily describe it in one sentence, we want to prove that this technology works and that Einstein’s general relativity is correct, also in this strongly dynamic regimes and so on. Now we know this. We know black holes are there, we know how the theory works, we know that the techniques we’ve used are the right ones, so now the next questions are just more diverse. Like if we see this effect, there’s still this effect, are there so many black holes or fewer black holes? And suddenly one has to position oneself new because everyone now has to define which of these questions do I want to answer. It’s not all of us... Previously all of us were running in the same direction. Now we are still going roughly in the same direction but everyone is covering slightly different aspects of it.” (Scientist III)

“Everything since then has been a little bit different. We know that we’ll have a future but we don’t know exactly what that future is gonna hold. There’s a little bit of uncertainty about how much we are going to gain from more events, like we need to start becoming a different kind of astronomer, not just making it first, but also building catalogues and combining data and (...) The focus is very different.” (Scientist II)

A PhD student working in one of the gravitational wave research groups experienced a change in focus after the first detections. The student has already been a member of the LIGO and has done research on neutron star merger waveform modeling. A neutron star merger with electromagnetic counterpart was detected in August 2017, just four months before the interview took place, when the student had already been a member of the research group in Hannover. However, since at the time of the student’s arrival in the Hannover group there had been more data from the black hole waveforms, the student chose the topic accordingly, depending on the need of the group and the newly provided data resources:

“My topic for the PhD is different than for the Master and... Why choose that? Because... now it depends on the need of the group I think. (...) Now we see that we have more binary black holes than neutron star (mergers) and so a lot of things have to be done, of course, for both, but for example, because binary black holes there’s more,

we need more faster waveform to generate, we need to fix that and we need more accurate waveform so that's what I am doing now.”

So what seems to be the thing to prioritize now is what best aligns with the newly acquired data, or what extracts most from the newly acquired data. For the PhD student it means moving from neutron star waveforms to black hole waveforms, for the continuous wave researcher maybe also moving somewhere where data and other resources are better used, and for the black hole researcher it means combining data from multiple black hole events. It is not the case that every research direction is open in the same way, since some detections are more foreseeable than others (such that louder events will be detected first), while some are (or have been) completely unknown or unpredictable (for example, will it be the black hole merger, the neutron star merger, or a very loud burst source such as a supernova). Since black hole binaries are being regularly detected, their fruitfulness lies not so much in the individual detections anymore but in collecting and comparing data in order to learn about their population in the universe and the range of their size, spin frequencies, or some other parameter of interest.

“Yeah, I mean there are very practical discussions now. For instance, if we see another binary black hole, is this even worth publishing, this is old paper, you know, two years ago this was, you know, a Nobel prize!” (Scientist III)

A particular case of priority, however, is urgency, and we can often find problems that are referred to as being urgent in the gravitational wave research publications. For example:

“We want to conclude our work with some remarks about the possibly most urgent question of the waveform model community. How well are we prepared for the upcoming era of advanced detectors and what needs to be done to improve the science output of gravitational wave detections?” (Ohme, 2012, 112)

But where does the urgency come from in the gravitational wave research community? The answer is twofold: it is on the one hand social, extra scientific, and on the other hand, it is epistemic, internal to the scientific enterprise. The epistemic urgency comes from the necessity to synchronize different aspects of the research, as emphasized in the following report:

“The urgency came from the fact that it's not just us theorists sitting there trying to figure something out and if we don't make it now we'll get it later, it was the fact that there is (an) experiment that relies on us, and when the experiment provides data, you know, we need to be ready and prioritization is the right word, you need to make sure that we do

our research so that we answer the most important questions first. We can't just follow, you know, every little idea that we have. We have to always keep the big picture in mind which to a large degree is set by the experimental data that we will expect in the next three years or at that point, you know. (...) And it's a question for everyone's own career of course, if I now research the right thing that is important, in two years I will become more useful, my career will become more prominent and you know, I can feed my family. But it's also, it's urgent to some extent for the entire community of gravitational wave physicists because we are in that niche and if we weren't at that time, if there would be no detections for another ten, fifteen years, the entire field might have died. So, yeah, it was urgent, in some time field to do the right thing." (Scientist III)

One source of urgency is extra scientific, social – if nothing is detected, the funds are withdrawn and the field dies. The other source of urgency is internal, epistemic, and has to do with the relation between theory and experiment. They should not come apart since they rely on each other. Gravitational wave modeling has proved to be well suited for analyzing the experimental results obtained so far, especially for the black hole binaries, but the theoretical and the technological aspects of gravitational wave research might diverge at some future point in time, as our scientist makes clear:

"It will be interesting to see whether our observational capabilities overtake our theoretical understanding at some point. So far, they seem to be going side by side, maybe our theoretical modeling is even a little bit ahead as we hope it is, but it could well be, I mean... the technology is advancing so fast, it could well be that in five years from now we have observations where we are not quite sure whether we understand properly what we have seen which will again turn the tides a little bit." (Scientist III)

In order to understand the data that has been commissioned by the interferometer, modeling has to improve alongside the experimentation. Louder events that have been detected allow for such refinement towards more matches with the signals, but also towards understanding diverse details of the signals, which is a way to extract more information about their source objects. Synchronization of the technological and the theoretical advances is crucial for the achievement of epistemic goals.

In addition to synchronizing these different aspects, a lot of time is spent now on reviewing work from competing research groups. Comparing and crosschecking results with other research groups has been necessary to establish the validity of observational claims, but after

the first detection and the consequent dispersion of research interests, this practice also becomes more time consuming and unfocused. Competitiveness among researchers is more pronounced:

“Times have become more aggressive now. (...) You know, elbows are having to do with it.” (Scientist III)

New goals have to be set in the new circumstances:

“We have to find, we have to clearly define new goals now, and you need goals for every research group and we have to... with this extra competition now, of course, there are these outsiders who haven’t spent their career on gravitational waves, but they don’t have to. Now it seems to be an interesting way to go about the universe... they want to come. So we all have to defend our position or redefine our position now, which in the sense is very stressful too, and again there’s some time limit that is very hard to define because it’s now sort of science-society time frame that you are following. If I now sit back and relax too much, you know, “I’m part of the detection, wonderful”, then I’ll find out in three years’ time: no one cares about this anymore. You’ve shown your results, and other people might be overtaking now.” (Scientist III)

Let me summarize what has been said so far. In the first chapter, I have introduced LIGO and LSC, i.e. gravitational wave research via laser interferometry. In the second chapter, I have first challenged the view of gravitational wave physics as a case of slow science for at least two reasons: the first reason is that the starting and ending points are unclear, and the second reason is that there is no independent, absolute, or non-relative way of measuring the speed of achieving research goals, since we cannot compare different kinds of complexities involved in different research fields. After that I have showed how setting of a particular epistemic goal, namely detecting gravitational wave signals from binary inspirals and burst sources, has been conditioned upon the possibility of cross correlating signals from two detectors. In the case of binary inspiral sources, it has also depended on the theoretical understanding of the waveforms produced by such collisions. The decision to focus on inspiral and burst sources by building two interferometers was a decision to set an epistemic goal that can actually be achieved in a reasonable time frame, which also means reasonable with regards to other resources: money, time, scientists’ interest, computational power. Moreover, it was an epistemically superior decision because the detectors running in coincidence make it possible to discern the gravitational wave signal from the surrounding noise by cross correlating two devices. The key

feature that allowed for a successful detection was the fact that the detectors were commissioning data in synchronicity, simultaneously.

In the third chapter, I have discussed prioritization and urgency in gravitational wave physics. I have shown how prioritization in terms of goal setting operates on different levels of the gravitational wave research: the individual level, the research group level, and the level of the research field. Individually, in the case of a researcher who is considering changing his research focus because there are less foreseeable results in his field than in another field. On the group level, in the case of a PhD student who is changing focus towards what the group needs the most. The group needs analyzing and refining waveforms for which adequate data is available, and not working on less frequent or yet undetected waveforms. Prioritization on the level of the research field, finally, seems to be an open question. There used to be a common goal, but now there are many, and it is not clear on what to focus. The groups have to position themselves anew. More importantly, new goals are being set based on the kind and the rate of the new advances made. With each new event, the groups restructure and redistribute resources in order to analyze the event as thoroughly as possible, sometimes at the expense of neglecting momentarily not as fruitful research directions.

Furthermore, I have pointed to certain epistemic (internal), as well as social (external) sources of urgency in gravitational wave physics. Externally, if there would be no measurable advances, the funds would have been withdrawn, researchers would lose their jobs, and the field would have died. Internally, with regards to the goal of attaining knowledge, the theory has to be ready for the experimental data, and vice versa – if the laser interferometers had not been built, the theory would have not been proven. The central aspect of epistemic urgency and prioritization is the need to achieve synchronicity and simultaneity – between the two detectors running in coincidence, as well as between the theory and the experiment.

4. Time and methods

4.1 Matched filtering

We will now move from the goal setting domain to the design of more specific research methods. Methodology has already been tackled in the explanation of the role for two interferometers in searches for gravitational waves from inspiral and burst sources. The promises of cross-correlating two devices established the initial research goal, i.e. the focus on

these specific sources of gravitational waves. In addition to cross-correlation, the possibility to theoretically model the inspiral binaries has been the advantage that enabled the first detection. It has already been said that the theorists model waveform templates and then the interferometric data is compared to the data in the waveform template bank through the method called matched filtering. The filter maximizes the signal with respect to the noise around it. The binary coalescence searches target gravitational wave emission from binary systems with individual masses from 1 to 99 solar masses, with total mass less than 100 solar masses. This is the range of most probable and detectable events and approximately 250,000 template waveforms are used to cover this parameter space (Abbott et al. 2016a, 7).

The challenge is to develop methods to solve the equations from which the gravitational wave researchers learn how the signal of two colliding black holes or neutron stars looks like. What has to be solved is a nonlinear system of partial differential equations, which cannot be done precisely. Different parts of the wave are modeled with the help of different methods. A gravitational wave signal from a binary has three parts originating in three stages of the collision event: an inspiral, a merger, and a ringdown. Inspiral is the stage at which two astronomic objects are orbiting each other at ever greater speed, until they merge in the merger phase when most of the gravitational wave emission is released. Ringdown is their “settling”, when a new black hole or a neutron star is created from the colliding two.

One of the methods used is post-Newtonian analytic approximation, which is especially useful in modeling the inspiral stage of the waveform. It is based on simplifications that are very successful for objects which are not too fast. However, the simplifications break down when the astronomic objects are too close to each other because they become very fast and highly relativistic, so the solutions arrived at through simplifications are not accurate any more. At this point another approach is needed – numerical relativity, which models the very last inspiral orbits of the collision and the merger. It can solve the full size equations so there are not many simplifications, but it takes a lot of computation time. The gravitational wave researchers need to combine these two methods in order to find out how the signal of, for example, two black holes merging into one looks like. Post-Newtonian analytical approach is accurate only when the black holes are far apart in the early stage of the inspiral and in low frequencies, but it is fast and it can calculate thousands of signals in a second, which is what is ideally wanted. The numerical relativity is very accurate but is extremely slow so it can only be used to calculate the last few orbits of the binary inspiral which require high accuracy.

Both post-Newtonian approximation and numerical relativity predict the gravitational wave signal, but the question is whether both methods will give a consistent representation of the last seconds of the inspiral, right before the merger event. If they do not, either one, or the other, or even both of them are not correct. Since each of them is measuring a different part of the waveform, what is needed is an overlapping domain where both approaches bring about approximately the same result (see Ohme 2012, 35-74). Still, even a high degree of agreement in the overlapping region of the two approaches “does not necessarily diminish the uncertainty of many choices that enter the modeling of (up to thousands of) GW cycles in the inspiral waveform” (Ohme 2012, 111). A balance has to be struck between optimally using a fast but less accurate method and a slow but highly accurate one, since they will have to converge around the same result at some point. The challenge is to get to that point as soon and as efficiently as possible.

Furthermore, prior to applying the model to a set of physical parameters, it has to be known where in the parameter space the models have been constructed. A parameter space is the space of all possible combinations of values for all the different parameters contained in a particular model. An isolated black hole is described only by its mass and spin, while two merging black holes are described by eight intrinsic parameters: masses m_1 and m_2 , spins S_1 and S_2 , location, orientation, time, and phase of coalescence (Abbott et al. 2016b, 1-2). The first observed black hole binary, for example, included black holes of 36 and 29 solar masses, resulting in a black hole of 62 solar masses, which means that the predicted waveform was in that part of the parameter space with respect to mass. As we have seen, the template bank includes 250,000 models for masses ranging from 1 to 99 solar masses. The researchers have to make decisions about the usage of their limited resources when it comes to modeling waveforms and covering particular parts of the parameter space. With any given amount of resources, should many short simulations be generated or a few very long ones? If less time is spent but more simulations made, it is possible to establish a parameter space, which means that it is possible to soon have more different signals modeled. Contrary to that, if more time is spent on each simulation, then only few waveforms can be simulated.

These considerations, again, reflect a concern about resources, primarily the resources of time and computational power. We can see how an overarching epistemic goal breaks down to more particular goals which can be achieved by certain accessible methods. Although what is sought for is an accurately represented waveform, in certain domains accuracy will be less important. For example, in order to have as many simulations as possible, the faster approach with short

simulations will be sufficient. Also, in the inspiral stage of the binary merger, simplifications will model the waveform correctly enough. Choices with regards to what is specifically wanted and how to get there are inevitably made.

“Which post Newtonian analytical approximation formulation should be employed?

What physical parameters in post Newtonian analytical approximation and numerical relativity are consistent with the other framework?

Which numerical relativity resolution, extraction formalism etc. is sufficient?

How long do the numerical relativity waveforms have to be?

What is the appropriate way to match analytical and numerical data?” (Ohme 2012, 75)

Among the choices are also those about which particular parts of the signal to focus on and which to ignore, at least until the limitations of the models are overcome.

“And often, when we look very carefully, we see that many, many parts are recovered correctly, and some parts are covered less accurately but then we don’t care, we make the choice to say: but this is fine, we can live with the fact that certain effects just don’t measure very well.” (Scientist III)

The modeling methods are geared towards overcoming limitations of one kind or the other, while pursuing very clear epistemic goals. I will now let a gravitational wave researcher reconstruct the process in his own words.

“Ideally, we have a little study, a well-controlled study where we find out how inaccurate can I get away with, right, (...) because this will tell me, this will give me the optimal usage of my limited computer resources. And what optimal means, this depends on how good is my instrument, how loud are the signals that I am going to expect, but this I can... You know, from talking to the experimentalists they can tell me: OK, in the next five years we’ll have this sort of instrument and this instrument will give you this sort of data, and the data itself comes with some accuracy, you know. So, so basically what I have to do is I have to make sure that my modeling is as accurate as or a little bit more accurate than the data. (...) So we’ll try to do a little study, and then, part of it is just generalized result of this little study, because the little study tells me: I need ten waveforms to cover this part of the parameter space. And then I calculate this numerical wave from this as good as I can and then I combine them with the more or less accurate

analytical models and then we go step by step, and see where it gets. And then we compare it... Often we now compare it with the competitor's group actually. So they will bring out their models and they will argue that their model is better and we will argue that our model is good. And then we compare. The parts that agree with more confidence, then they are probably right, and the parts that disagree... But at least one of the models or probably both of them are probably accurate enough so we'll invest more time in those." (Scientist III)

A controlled study is done in order to estimate the error between the results of the post-Newtonian approximation when compared to numerical relativity. The error estimation is important because it optimizes computer resources. What is optimal also depends on what the experiment is able to commission: it is not useful to make predictions of a waveform whose frequency is beyond the sensitivity level of current devices. Interferometers get "dirty" data, buried in the noise, with certain parts of the waveform not measuring well. The predictions therefore have to be more accurate than such data, they have to account for what is not being "heard", i.e. measured well. There can be many waveforms modeled by the "fast" post-Newtonian analytical approximation, so they start from a particular numerical simulation generated waveform and see with which waveform generated by the other method it aligns best. After that, they compare the completed waveform to the competing group's model. The goal is to gradually cover more and more of the parameter space, and eventually to find the best match with the candidate signal.

A particularly interesting part of the researcher's report is "finding out how inaccurate can I get away with". Though accuracy is highly desirable, it comes at a cost of more computation time. Most of the time, it is simply not completely achievable ("data itself comes with *some* accuracy", emphasis mine). So what researchers do is try to find some room, some space of acceptable error that will still lead to useful predictions. In this particular case, it means how many spins will be modeled with the analytical method and at which point near the merger will the numerical relativity take on. The results of the two methods have to be consistent, so consistency might be the criterion for being accurate enough. Still, consistency itself comes in degrees so the problem is only further relegated to what it means to be consistent enough. As we have seen, sometimes the choice is simply to ignore a part that is not measuring well. Uncertainty inevitably permeates the modeling process.

“The important conclusion we shall draw from this is that none of the complete IMR³⁷ waveforms is based on an unambiguous construction, and the spread of possible results that different reasonable choices yield is a measure of the uncertainty within the modeling process.” (Ohme 2012, 75)

What I aim to convey is the amount of decision making that is employed in order to overcome various limitations, including those of time. Specifically, temporal limitations are strongly intertwined with the limitations of cognition and computation. The performance of cognitive or computational systems is measured by their outputs as delivered in a particular amount of time, such as a second or an hour, in the same way in which a student gets a grade in a written exam based on solving the exam in a certain time slot, and not solving it in principle.

I will call this a math class example. A different skill is being demonstrated by solving a homework that is taken home and written up until the next day, as compared to the skill that is demonstrated in particular constrained circumstances, such as those posed by a written exam format. The former is designed to enhance ability, while the latter is supposed to measure efficiency. This is a highly simplistic account of what is going on in educational settings, but it maps well onto the waveform modeling issues that we are discussing. Both methods for calculating waveforms possess certain ability or value – to be very accurate, or to be very fast and cover a broad range of possibilities – as well as certain efficiency in exercising one or the other value. While post-Newtonian approximation demonstrates the ability to be fast under the constraint of being sufficiently accurate, numerical relativity is able to be very accurate, but only when unconstrained. Since the unconstrained is only rarely optimal, it will be utilized primarily as an error estimation. Its function as a control for the accuracy of the analytical method will enable the researchers to use the available resources as efficiently as possible. In this case it means maximizing the usage of minimal computing time by the analytical method, while minimizing the usage of maximal time by the numerical method.

4.2 Machine learning

Validation of a waveform model is done in three steps: first, a comparison of the analytical and the numerical method is done, second, the resulting waveform is compared with the competing group’s waveform, and finally, the model is validated by the matching data and gets refined as

³⁷ Inspiral-merger-ringdown

more matches are found. Another way to make searches more efficient is by using machine learning methods:

“So the other part is trying to find new ways maybe from mathematics or so to do this stuff that we are doing more efficient, because some of these combinations of numerical data and simpler methods... In some extent it’s quite naïve, you know, we are physicists, we just want to solve the problem, but mathematicians have thought about this on a much more abstract level so we are trying to understand what they can offer us which is not easy because they speak a very different language, but the hope is that we find new mathematical tools that allow to make this combination even more efficient, so basically to train the computer that can then tell us: this is important, this is not important, don’t do this, which... a lot of the things we are doing now is partly intuition and partly sort of humans looking at it and asking questions but I think... A fair amount of this can be solved by computer, machine learning, and all this kind of stuff.” (Scientist III)

The account of the scientist highlights the differences between the deductive mathematical method and the empirical method of physics. The physicist calls their own approach naïve, which is further clarified by appeals to problem solving, intuition, and “humans asking questions”. Contrary to the empirical method, mathematics is deductive, and apparently more efficient in reaching true conclusions. Apparently, because not any truth is the object of scientific interest, but rather truths about the empirical world. The empirical method in LIGO is, however, indispensable from mathematics from the start – what the waveform theorists are doing is solving nonlinear equations. These solutions are mathematical, only applied to a set of physical parameters. The simulated models are called hybrid waveforms because of the combination of analytic and numerical methods that have created them (Ohme 2012, 45-53). When it comes to the application of machine learning to waveform modeling, the goal is to “train the computer” to make the kind of decisions humans usually make (“this is important, this is not important”). Decision making about relevance, however, is empirically informed and involves a possibility of error, i.e. it is inductive.

Two recent articles explore the possibilities of deep learning algorithms, a subset of machine learning, in substituting matched filtering in searches for gravitational waves from binary transients (George and Huerta 2018; Gabbard et al. 2018). They use simulated signals from the waveform template library to train the computer to recognize a gravitational waveform in the noise. The role of machine learning in substituting computationally costly methods might turn

out to be invaluable for the future development of the field, since it can increase computing speed. The results of machine learning tests as compared to matched filtering are fairly optimistic, showing a close match in sensitivity and allowing for improvement by increasing the number of training samples (Gabbard et al. 2018, 5). By relegating the searches for matches to computers, researchers can redistribute cognitive and computational resources in a more optimal way and work on a broader range of questions at the same time. This makes them epistemically better off in a given amount of time. In other words, the epistemic goals are met and time is saved.

4.3 Searches for continuous waves

So far, the focus has been on the binary inspiral modeling, primarily the black hole binary. Another kind of challenge that gravitational researchers have to overcome is how to use resources and time available on computers to detect things that have not been seen yet, like the continuous waves from spinning neutron stars. The black hole and neutron star mergers that have been detected so far are transient events, very fast and very loud, while the sources of continuous signals are fairly quiet, but they go on for a long time. Since they last for years it is possible to see them by using a large amount of data, but this is very computationally intensive and one needs to be able to find the signal in the noise.

One of the ways to do this is by using a method called cross-correlation resampling. It is used for searching a long signal for a shorter, known feature. Resampling techniques can accelerate the cross-correlation method. A continuous signal is expected to match (correlate) with a pattern over many sequences of data sets from a single detector, as opposed to a correlation between two or more sites for a single extraordinary event, as it is the case in binary and burst searches. The computer code needs to be accelerated to make new searches possible, since a continuous signal has to be caught exactly over billions of cycles over the course of a year or millions of cycles during the course of days. A vast amount of different possibilities has to be tested on the computer, which takes time. The best way to look at these weak signals is to “chop” the data into shorter bits and look for a signal which matches with the shorter bits, and then combine them all together. Ideally, one would look at the data at once, but that is impossible.

“It would take the life of the universe to do it. By breaking it up into small chunks it’s not as sensitive but it gets more sensitivity for the same computing power.” (Scientist II)

This is also a way to enable a synchronized search: since it is impossible to look for a weak and long signal at once, data sequences are simultaneously analyzed for the same known feature. In this case, one detector is enough since the replication work is not done by the other interferometer, but by an identical pattern occurring in multiple data sets commissioned by a single instrument over a long period of time. A version of replication is necessary in order to cancel out confounding. The replication of a continuous signal is supposed to happen in time: the same pattern will occur in different sequences of a single long signal. Cross-correlation is thus a form of replication, achieved either by spatially distributed instruments detecting a loud and short signal at the same time, or by searching for temporally distinct sequences of a weak and long signal in the data acquired by a single instrument (or more of them in some cases). The searches for continuous waves have still not resulted in success.

“This is always something we try to keep in mind in science... is what is the room for growth. Like, if somebody gave me ten times the computing resources, for example, I would immediately know how I could make my search better. I could just increase the amount of... the size of the segments I look for and I could be more sensitive. (Scientist II)

The increased size of data sets enables searches in longer sequences which allow more sensitivity because a continuous signal is supposed to propagate in time with the same amplitude and frequency. The longer the data set, given the adequate computational resources, the easier it will be to discern a continuous pattern of a certain frequency and amplitude range. The more computational power, the longer the segments and better the search. We can recognize the need for enhanced and accelerated algorithmic searches and machine learning methods in matching the filters with the data, since, again, it is crucial to optimize computing resources.

In order to make tests possible by adding more computing power, a volunteer distributed computing project called Einstein@Home was set up in 2005. Einstein@Home uses the idle time on people’s computers all over the world to search for weak signals from spinning neutron stars, using data from the LIGO detectors, the Arecibo radio telescope, and the Fermi gamma-

ray satellite. It is a way to get ten times as many resources as, for example, the ATLAS supercomputer in Hannover has.

“So you might give a particular user a set of possibilities to examine and say: “Could you look at the data and see whether there might be a signal from a spin frequency of 100 Hz to 105 Hz and maybe in this portion of the sky?” And then their computer will test those possibilities and report back: OK, we saw that something might be there with this significance or this statistic, it’s this thing there; and then Einstein@Home pulls back the results from everybody and concatenates them, joins them together and looks for a signal.” (Scientist II)

Einstein@Home volunteers have already discovered about fifty new neutron stars. The community web pages of the project are a site of exchange among the laypeople involved. Here is a post by an enthusiastic contributor which exemplifies the kind of contribution by the public that is involved in Einstein@Home:

“13 new gamma-ray pulsars! This news inspired me to resurrect an ancient server and update it with a couple of \$10 CPUs (E5540) and inexpensive RAM, and last week I added a modern GPU and now have a reasonably powerful box dedicated full-time to Einstein@Home. And I just got a friend to join, too. Now if I can just get the GPU overclocked a bit...” (A comment posted by user Tachyon on March 20, 2017 on Knispel 2017)

A distributed computing network increases computing power and distributes targeted searches over a broad range of parameters. It is a way to enable simultaneous searches for different variables with the aim to gradually cover more and more of the parameter space. In a very peculiar way, it is also an instance of inclusive, citizen science, since lay people are invited to contribute, although none of their expertise is really needed, because what is needed is the idle time of their computers. Nonetheless, it is a platform for non-expert involvement and thus epistemically and politically important, in line with the attempts to democratize science. But most importantly, it is a way to make searches for continuous signals more efficient.

In the searches for continuous waves, we can again notice a trade-off between longer segments with greater sensitivity and shorter segments with less sensitivity, but which allow for speed ups by the cross-correlation resampling method. Again, the searches are enabled by the availability of resources, primarily computing power and time. An especially prominent feature of continuous waves is their long duration which poses specific problems for their successful

detection, since time series have to be integrated over a very long period of time, which is especially computationally intensive. We have also seen a possible solution to this challenge – a distributed computing network which optimizes computing power and time resources.

5. Time and the evaluation and communication of results

5.1 Introduction

Alongside the increasing number of detections and technological advances, LIGO has made great efforts in public outreach and science popularization since the first detection. Public outreach is a stated objective in the official documents and it is made visible on the official web pages where various educational, media, and outreach tools such as the LIGO Magazine are available. For example, it is possible to get real time alerts when a new candidate for a detection is found. The first detection, however, was kept secret for five months. This was necessary in order to make sure that the event was real because the epistemic input and the scientific impact of a possible detection have been huge, and the expectations on the part of the general public and the scientific community that might make use of the results have been high. It was thus important to avoid false positives as much as possible, of which I will say more in the next section and in part III.

Sociologist of science Harry Collins has devoted his book *Gravity's Kiss* (2017) to the first detection of gravitational waves, because this event was not only a scientific finding, but also a carefully negotiated social process. A detection does not happen in a self-evident way, but it has to be thoroughly analyzed so that the possibilities of confounding, systematic errors, instrument noise, human errors, and blind injections that were a part of training for the actual detection can be canceled out, and that the detection can be announced with a high degree of confidence. Moreover, many non-epistemic and pragmatic decisions have to be made in communicating the discovery to the public.

Blind injections are fake signals that mimic a real signal so that researchers can learn from a simulated detection in order to gain expertise for the real event and thus anticipate possible problems. In the two cases of blind injections in 2007 and 2010, the majority of researchers did not know that the signal was not real and have learned about it fairly late in the process of

analysis.³⁸ Many of them thus had a reason to initially suspect the actual detection when the candidate event happened because the match was thought to be “too good to be true”.

“And when we saw the signal, actually at first I thought, that’s kind of funny, now this can’t be real because it’s too good. It matches too closely to what we’ve expected... (...). So, there was a phase where I was wondering, this can’t be true, this is too good to be true.” (Scientist III)

A blind injection involves a coordinated manipulation of at least two interferometric data sets, so it would be extremely hard to achieve the degree of precision with which the actual detection fitted the predictions. LIGO management was from early on insistent that the September 2015 event had not been a blind injection, but a protocol was nonetheless set to cancel out a malicious injection by somebody from the team. For the same reason of extreme degree of expertise and coordination that would have to be at work in order to perform such a deceit in a successful way, a malicious injection was soon out of the question.

The confidence in reaching project objectives was often low and the trust among the members of the group has sometimes been shaken due to high stakes (in terms of epistemic gains and reputation) involved in such a big project comprised of dispersed groups. This was especially so in the absence of real detections coupled with the experience of blind injections, in addition to many instances of opposition to the project from early on. There was thus no room for mistakes when the detection candidate event was recognized in the data, as it was stated in a previous quote on p. 29 by González and Cavaglià (2016). Many decisions had to be made through a negotiation process including over a thousand LIGO members: who will be responsible for the draft of the detection paper, how widely outside LIGO and LSC can the information about the event be revealed prior to publication, how will the event be named, will it be called an “observation” or a “detection”, how long will the draft of the detection paper be open to revisions, and who will be included to revise it?³⁹

Some decisions had been made by the management while some had been open for a broad discussion among the members, though the discussion had to be strictly temporally constrained in order to move the process efficiently through the evaluation and publication steps. These decisions had to be made quickly in the months following the detection and preceding the

³⁸ Collins wrote about the two blind injection events in the book *Gravity’s Ghost and Big Dog* (2011, 2013).

³⁹ See Collins (2017).

detection paper and media release. The researchers had to be very certain in the event, but also very fast in revealing it to the world, since the detection was expected with great interest due to the project's epistemic significance, but also due to substantive investment without tangible results so far.

5.2 Non-epistemic and pragmatic values in the communication of results

A notable non-epistemic revision had already taken place in the course of the LIGO timeline with respect to the terminology used in specifying the very aim of the project. Gravitational waves had until the last few decades been called “gravitational radiation”, but during the Cold War a reference to radiation was thought to be detrimental for the public reception of the project. “Radiation” had with time been replaced with “waves” in order to avoid unwanted associations with nuclear radiation. Furthermore, the first observed (or detected) event was eventually called “observation” in the title of the paper, while the term “detection” was included in the body text. Astronomic observatories had been unsatisfied with LIGO's appropriation of the word “observatory” from the start of the project, because the search for gravitational waves was not a part of mainstream astronomy. It was rather in the realm of fundamental physics. It had thus been important to state clearly in the title of the paper that what had been achieved is an observation, because LIGO is an observatory. “Detection” is nonetheless an even more widely used description of the same scientific event. The fact that it was the first direct observation/detection was left out of the title so that success would not be undermined by indicating that the waves had already been indirectly observed.

These are examples of decisions that reflect non-epistemic, but also pragmatic values in the process of communicating results. They are non-epistemic because they are not conducive of truth, and they are pragmatic because they are instrumental to the aim of communicating results to the public efficiently, which contributes to the overall knowledge distribution. Thus, they can also be understood as epistemic. For example, in Alvin Goldman's social veritistic epistemology framework, knowledge is increased through communicative or testimonial practices which enable the recipients of knowledge to ultimately accept true beliefs.

“Communication is an efficient mode of increasing knowledge because information transmission is typically easier, quicker, and less costly than fresh discovery.” (Goldman 1999, 103)

The more people come to acquire knowledge, the overall epistemic gain is greater, and a social epistemic goal is met – an increase in knowledge distribution which can lead to an increase in new knowledge acquisition. Under this framework, the efficient communication of new findings is an epistemic aim, and thereby decisions that are instrumental to this aim are epistemic decisions too. This leads to an apparent conflict since it was said that the examples from this introduction reflect non-epistemic values. For example, the decision to talk about “radiation” rather than “waves” does not contribute to the goal of detecting the very phenomenon at stake, although it increases public support for the project. However, if public support increases the likelihood of engaging in scientific practices that will eventually achieve the detection, then it ultimately plays an epistemic role. Daniel Steel (2010) distinguishes between intrinsic and extrinsic epistemic values. The distinction depends on whether a certain value or practice is an indicator of truth, or it only supports or contributes to truth but does not indicate it, as it is the case with choices about terminology. In part III I will present Steel’s account in more detail and evaluate it. I will ultimately supplement his extrinsic vs. intrinsic epistemic value framework with a more explicitly pragmatic understanding of epistemic values.

The claim that I will advance is that pragmatic considerations such as those about the speed of practices in generating results are inseparable from epistemic values in the pursuit of epistemic aims of research. Pragmatic dimension is inherent to epistemic values, since limitations of resources are an inherent dimension of knowledge pursuit, in the same way as the distribution of truths is an inherent dimension of knowledge in the social world. Decisions related to the speed of practices sometimes overlap with non-epistemic considerations, but they carry a more substantial epistemic role. Epistemic aims are pursued in constrained circumstances of limited resources of time, cognition and computational power, so decisions that serve the requirements of these limitations are thus instrumental to achieving epistemic aims.

Some of the examples are: the decision to focus on binary inspirals, to build two interferometers, to combine methods in waveform modeling, to use machine learning methods, or to set up a distributed computing network. On the level of goal setting, pragmatically informed epistemic aims determine the research focus on different organizational levels. On the level of communication, pragmatically informed epistemic aims have brought about a transition from gravitational radiation to gravitational waves. Many of these decisions are especially dedicated to ensuring that the resources are used efficiently. This inevitably involves a temporal dimension since efficient means successful and timely. Both the category of success and the category of timeliness are dependent on the particularities of the case involved, i.e. they are

highly contextual. Decisions about how to best serve particular epistemic aims are inevitably permeated with pragmatic considerations that address various limitations, among them those of time. In a word, pragmatic is epistemic.

I will now turn the focus on clearly temporal aspects in that last phase of research, prior to the announcement of the detection. As already stated, the evaluation and communication had to be done quickly yet thoroughly, in order to yield great certainty. Some of the limitations to this effort came from the fact that decisions had to be negotiated among many people involved in the project. Some limitations came from the fact that the longer they waited, the less likely it was that the detection will be kept secret. Leaking of information might have compromised the research since it would increase the expectations and the pressure to confirm a detection, and a false positive in this case would be damaging the overall trust in the project. The public and the broader scientific community were therefore not made aware of the candidate event. Also, the more the time would pass, the more suspicious the scientific community would be as to why the outcomes of the second observation run are not communicated and made available for usage by other research groups. Hiding of such important information might also compromise the trust in the project, in the same way as it would be damaging to prematurely release an information that would turn out to be false.

On the other hand, a relatively short time to publication inevitably involves trade-offs, since not every aspect of the detection could have been analyzed by the time of the announcement. The decision was thus made to first issue a detection paper with only the basics, and then subsequently publish more detailed analyses, rather than publish everything together that would then had to be delayed for a longer time. It was also a matter of negotiation what exactly are the basics, since there was disagreement with regards to what should be included in the detection paper. Luckily, the first detection was estimated a remarkably low false alarm rate, a high signal to noise ratio, and a statistical significance level of 5.1 sigma, which is above the established golden standard for physics, 5 sigma – a p value of 3×10^{-7} , or about 1 in 3.5 million chance of a mistake (Lyons 2013; Staley 2017, 48-53). Collins argues in his 2017 book that if some of the subsequent events happened to be the first event, LIGO would have a harder time in convincing the public and the peers that it was a real detection. With the credibility gained by the first event, the subsequent ones were relieved from a great deal of pressure to “overdo” the evaluation. They could also focus on different things after the ice has been broken with the first event.

Furthermore, since LIGO is not a project of immediate application and societal impact, all of the decisions in the phase of evaluation and communication were related to the publication paper, since this is the format through which a credible outcome of this kind of research is communicated. Scientific publication is a medium that makes the result available for evaluation and usage by the scientific community. In contrast, a medical therapy gets used already during the process of evaluation, namely in clinical trials (recall Ebola *ca suffit!*), while a scientific paper is written up last, after the preliminary results of the trial have been made available. It is possible that a therapy is in usage before the paper is written, since the primary target group are patients and not scientific peers, and the primary goal is often application and only then knowledge.

In gravitational wave physics, it is the peers that have to be convinced in the result, since this is the only way it can be communicated to the public that otherwise does not have a sufficient degree of expertise to assess the significance of the finding. In more obviously applied research contexts this is remarkably different, since what is sought for is a product, an application, and not peer-reviewed knowledge. A patient will often subjectively feel that a new medical therapy has beneficial effects, or she will buy a newly designed piece of technology with an additional performance, while there is no change to trace in the world with a gravitational wave detection having taken place, and the evaluation and communication tool is (only) a high impact peer reviewed scientific journal. A press conference was held and a media release about the first detection was issued on the same day as the detection paper was published, in order to inform the public about the scientific event in a comprehensible and appropriate way at the same time it became available to the scientific community. This synchronization also had to be carefully coordinated in order to satisfy the requirements of both scientific rigor and technical detail for the peers, as well as of accessibility and clarity for the non-expert public.

Finally, inductive risk considerations come to focus in the evaluation of candidate events. Inductive risk is involved whenever hypotheses have to be accepted or rejected. In this case, when a candidate event has to be evaluated as a possible gravitational wave signal. Two mistakes are possible: that there is a gravitational wave signal, but is wrongly ruled out as one (false negative), or that there is no gravitational wave signal but it is wrongly recognized as one (false positive). Researchers have to make sure to avoid false positives as well as false negatives, but often they cannot do both at the same time, since there is a trade-off relation between the two. If one highly cares to avoid a false positive, one will have to be more tolerant of false negatives, and the other way around. Since LIGO is a long-term effort and the epistemic

stakes are high,⁴⁰ negatives of any kind are basically business as usual: no detection, like in the last decades. There has been a theoretical certainty that gravitational waves exist, but the sensitivity of the experiment had been inadequate for any detection until recently. A gradual change had been expected, but a negative result, including a possible false negative, has been an expected status quo for quite a long time. Interferometers pick up all kinds of noises and signals, and nobody cares to rush to announcements of detection when it is much more likely from previous experience that there is no gravitational wave signal than that there is a gravitational wave signal.

It has been a given that a detection should be achieved with great certainty if it is to count as real in the wider community of scientists and public. This has put significant pressure on avoiding false positives, since this type of error could have been a final punch for the project. Credibility would be lost if researchers would be more likely to tolerate false positives than false negatives. False positives were precisely the reason why prior attempts to detect gravitational waves had lost scientific credibility and financial support. When the scientific community had not been convinced that the sensitivity level was adequate for a detection, then the insistence on claims to the contrary only contributed that those who advocated them lost support for their work.⁴¹ It was therefore of utmost importance to avoid false positives because they would cause a greater epistemic loss as compared to false negatives. First of all, because that would run the risk of building further knowledge on a false fact. Second, because the project might lose funding, and then the insight that it is supposed to enable would not be achieved.

On the contrary, a false negative, a situation where there is a wave, but it is not detected, is pretty much what we have all the time, and it is tolerated to a greater degree, although ultimately far from wanted because it also leads to a loss of epistemic insight and resources. For example, it is known that continuous gravitational waves are buried in the noise somewhere in the data, but the searches have not yielded in success because the needed sensitivity is not yet reached. It is a long-term problem for the project if the signal turns out not to be found for a long time

⁴⁰ Non-epistemic stakes in this case refer to lost resources, but that does not have an immediate impact on anything, although it can decrease public support for similar projects. High non-epistemic stakes would be losses of lives, health or other value related to life, society, bio- and eco-systems, or economy. I doubt that anything like that is at stake in the case of LIGO.

⁴¹ This refers primarily to the case of Joe Webber and early resonant bars which is described in detail in Collins (2004).

still. However, it is not a problem immediately, and not even for some time yet. A false positive, on the other hand, is a false hope, a threat to the credibility of research efforts and trust in scientific integrity. It might compromise the project more directly. However, we can see a change after the first detection. It is much more costly to tolerate false negatives now, because the sensitivity of the interferometers has been proven adequate for detections, which has brought about a change in goals. Now it is not only important to detect something, it is much more important to keep detecting in order to, for example, learn about the distribution of various sources of gravitational waves. Precisely for this reason have methods such as machine learning and accelerated searches been advanced – the proof of concept has been acquired and the desired sensitivity and significance rate has been achieved. In the new circumstances it is becoming much more important not to tolerate false negatives, while an eventual false positive would not have as negative effects as it would have had before, had it occurred prior to the first detection.

Furthermore, acceptance of a hypothesis, in this case that a particular signal is a gravitational wave signal, is not only evaluated according to a standard that can guide exclusively researchers' belief about the hypothesis, it is rather evaluated according to a standard that can make others believe in the result. The five sigma significance level had to be reached in order to communicate the event to scientific peers and the public with confidence, not for the researchers themselves to believe that it is real (cf. Staley 2017). The initial preference for false negatives has thus not been chosen in order to influence researchers' beliefs about the truth. It has been chosen in order to enable action, i.e. to efficiently communicate results to others. The researchers have been convinced by the signal regardless of the ultimate sigma level. The aim of communicating results to the public, including the scientific community, is well served by incorporating considerations of how this should be done in a trustworthy, convincing and comprehensive way, and the only way is to raise the bar very high in terms of what should count as a positive result. Notably, such reflections about actions bring the inductive risk argument, which is usually tied to practical projects and societal stakes, in the heart of fundamental research. In the case of gravitational wave physics, the influence of non-epistemic values on research does not exert itself via concerns about people's health or safety in the case hypothesis is falsely accepted, but rather via a concern about the reputation of the project and the public trust related to it, if the gravitational wave is not in fact observed when it has been announced that it is.

6. Conclusion

In the last part of the thesis more will be said on inductive risk, epistemic aims, value trade-offs, and especially the pragmatic dimension of epistemic values. For now it suffices to be acquainted with the LIGO project, the first detection, and the current state of the field. The details of the case study will serve to show that pragmatic decisions are recognized in every stage of the research process: they influence goal setting, methodology, evaluation, and communication of results. These pragmatic decisions are an inseparable part of epistemic efforts and are conducive to the fulfillment of epistemic aims.

The focus has been on LIGO, an international collaborative project aimed at detecting gravitational waves by means of laser interferometry. After introducing scientific objectives of the LIGO Collaboration, its experimental setup, and theoretical underpinnings, my aim was to explore how considerations that address time constraints play out in different parts of the gravitational wave research. With respect to goal setting, we have seen how the initial goal to detect gravitational waves from binary inspiral sources was prioritized over other sources because of the possibility to single out candidate events by cross-correlating two or more interferometers located far apart. The decision to build two interferometers as opposed to one was informed by the epistemic goal to detect waves from these particular sources, since cross-correlation of two instruments was considered an epistemically superior method, because it would yield significant results more quickly and easily, thus convincing the funders and the public that laser interferometry is a scientific effort worth pursuing. Cross-correlation is a form of replication, in this case between two spatially separated experiments. An identical signal on two sites needs to have a common cause, hopefully a gravitational wave from a distant astronomic event. This enables a detection with a high significance rate which would be hard to achieve in the same time frame with only one detector, even of a much more refined sensitivity, since the problem of confounding could not be adequately cancelled out. Two interferometers have thus been an epistemically better option than one interferometer, among else, because they have been a quicker option. The two aspects, scientific and extra-scientific, are hard to disentangle in this case.

I have used researchers' reports, some acquired directly through interviews, while some taken from Collins' work and rich online resources that have documented first detections with great detail. This provided an insight into how considerations about time enter decisions about which lines of research to pursue and which goals to set for different levels of research organization:

the individual, group, and the whole research field level. I have shown how goals change with time and how resources, especially cognitive and computational, get redistributed with new prioritizations. I have also identified some internal and external sources of urgency in gravitational wave physics. Internal urgency pertains to efforts to maintain synchronicity between theoretical and experimental aspects of research. Prior to the detection, an external urgency, on the other hand, was to detect a gravitational wave signal in order to keep the funds. This was, however, also important internally – if nothing can be detected reasonably soon with a sufficient level of statistical significance, then the project does not deliver on epistemic promises.

Epistemic aims are pursued in the context of limited resources, so goals are set, methods designed, results evaluated, and finally, communicated, in a way which deals with these limitations most successfully. In this respect, speed of practices to generate results (to single out candidate events, to model waveforms, to match filters) is highly valuable, both epistemically and non-epistemically. Epistemically, as a way to reach true conclusions earlier on and thus create more chances for a further increase in knowledge. Non-epistemically, because of researcher's career process and maintaining funds, which is usually based on measurable advances on a politically relevant time-scale. LIGO has been considered a risky investment in this respect, spanning through several decades and with funds already waning prior to the detection. After the detection, the field got extensive funding again, with new research groups forming and new interferometers being approved and built.

Although I have started with the contrasting notions of 'internal' to science as opposed to 'external' to science for the sake of clarity, I will ultimately claim that the distinction is hard to withhold. The same will be the case in the next part, on translational medicine. In part II, we will see how external and internal aspects of science collapse into one in contemporary biomedical research, where societal stakes are more evidently involved. In the case of gravitational waves, the two aspects are seemingly more disparate, but we can nonetheless see how the internal/external distinction gets blurred. For example, extensive funding enables the fulfilment of epistemic aims, which are, on the other hand, not well served if nothing is detected relatively soon, since the interest and funds are thus lost. Also, if the needed significance level could not be achieved, despite possible epistemic advancement, nobody would be convinced of the detection, and again the epistemic aims would not be met. Goal setting and method design, as well as the evaluation and communication of results, are thus stages where the internal and external factors related to science more or less evidently overlap.

Trade-offs between different epistemic values, on the other hand, have been made most evident in the discussion of waveform modeling, characterized by great uncertainty. Many decisions have to be made, and the focus has been on trade-offs between a slow but very accurate method, and a fast but not as accurate one. I have shown in which ways researchers address the limitations of computing power and time, for example, by running a controlled study to learn “how inaccurate can one get away with”, or setting up a distributed computing network, or developing machine learning methods.

Finally, I have touched upon preferences for errors in the inductive risk considerations. False positives would be much more detrimental for a project like LIGO since they would resonate more dramatically with future research efforts and public expectations. It has been important to reach the necessary sigma level of significance in order to be able to announce the detection with confidence and open it for usage to other research groups. A preference for tolerating false negatives was especially prominent prior to the first detection, when the group yet had to establish its reputation with a finding that is supported by a high significance rate. A candidate event puts a lot of pressure on evaluation in terms of accuracy, efficiency, clarity, and timeliness. After the first detection, however, a preference for the type of errors also gradually changes.

A case study in gravitational wave physics has been motivated by what has been called an unfair comparison in the introduction of the thesis. The case study exemplifies fundamental research without immediate societal applications, open-ended in terms of both the timeline and the research goals. I have shown how considerations about time and speed nonetheless play a role in all stages of research, as well as on different levels of the research organization. The gravitational wave physics case study will be confronted with a case study in translational medicine, which is very directly socially relevant and whose applications are expected with great interest and need. Both case studies provide valuable material for a discussion of the pragmatic dimension of epistemic activities, a task that will be most directly taken up in the last part of the thesis. It will be argued that pragmatic considerations related to time as a resource play a significant role in scientific research, and that values that address temporal limitations of research, such as speed, should also be considered epistemic, since they have an invaluable role in the fulfilment of epistemic aims.

Part II

Case Study II: Translational Medicine

1. Introduction

An important aspect of current practice in medical science is the promise of speeding up of the research and development process from discovery to implementation in clinical practice, most explicitly present in translational medicine. Translational medicine is the use of basic and applied research to obtain new knowledge and develop skills, understanding, and innovative products which can result in implementable applications (Aronson 2017e). Biomedical research, including both the development of drugs and therapeutic practices, is especially dedicated to time-sensitive outcomes. In recent years translational medicine has become a prominent approach aimed at bridging the gaps between laboratory, clinics, guidelines, practices, prevention, and health politics. The idea is to make research more end-user centered, translating it into clinical practice and health policy in a more efficient, timely and patient-sensitive way, while maintaining rigor on the safety part.

The deliverances of medical research are expected with great interest and even impatience fueled by the amount of investment it receives, especially in the US. From 1950s on, the level of investment in health research and development has been increasing significantly, with a total of 171.8 billion dollars in the US for 2016 (Research America 2017, 3), while the number of new drugs approved annually hadn't been following the budget increase (Munos 2009). In the 1990s and early 2000s the number of new substances approved each year by the US Food and Drug Administration (FDA) as novel pharmaceuticals had been more or less stable and ranged between 20 and 30, with a peak in 1996 reaching 53 approvals (Mullard 2019). The change in 1996 is sometimes associated with the enactment of the Prescription Drug User Fee Act (PDUFA) four years earlier, which allowed FDA to collect fees from drug producers to fund the new drug approval process (Munos 2009, 959). PDUFA enabled FDA to acquire direct industry funding to hire more staff in order to review new products faster (Carpenter, Zucker, and Avorn 2008), but it has also raised concerns about the standards for drug approval (Horton 2001; Stegenga 2017). However, the numbers have soon dropped to their previous range, despite pharma money. Trends hadn't been good in the earlier decade either. Of 101 promising claims of new discoveries with clinical potential that were made in major basic science journals

between 1979 and 1983, only five resulted in interventions with licensed clinical use and only one had extensive clinical use by 2003 (Contopoulos-Ioannidis et al. 2003 and 2008). Recently, the number of FDA approvals hit a two decade high in 2017 with 46 novel medicines, followed by 59 approvals in 2018 (Mullard 2019). Nonetheless, many new approvals are in fact known drugs used for a new purpose or used to treat a new population of patients, for example children (FDA 2018a).

The reason for the recent approval rate increase might be in the strengthened academia-industry relation (Takebe, Imai, and Ono 2018; Robinson 2019) resulting at least partly from translational efforts (NIH 2014), but this yet has to be shown. It is beyond the scope of this thesis to definitely assess the success of the translational paradigm, partly because it is extremely hard to establish and measure the parameters that would be indicative of that success. Furthermore, having new drugs brought to the market does not necessarily mean that important health issues have been successfully addressed or that access to treatments has been improved. In fact, there is an increased awareness that this is not the case and that biomedical research and regulation is characterized by a number of shortcomings, hence failing to adequately live up to its social role (Bartfai and Lees 2013; Brown 2004, 2017; Reiss 2017; Stegenga 2017, 2018; de Melo-Martín and Intemann 2018, 96-115), although life expectancy has been increased and the quality of life significantly improved in the developed world in the course of the last century.

The shortcomings of biomedical research mostly dealt with in philosophy of science are related to the increasing commercialization of research, which has led to data often being misinterpreted, fabricated or ignored, in order to make a strong case for a potentially profitable product. This is especially so when it comes to pharmaceutical industry which has been responsible for two historic tragedies, the thalidomide case in the 1950s and 1960s, and the Vioxx case in the early 2000s. In both cases human lives were lost or severely damaged as a consequence of the usage of supposedly safe medications.⁴² Faulty research design coupled with commercial interest has been the cause of compromised patients' safety.

The general view is that pharmaceutical productivity has been going through a crisis for at least three decades (Munos 2009; Pammolli, Magazzini, and Riccaboni 2011; Taylor 2016). The advances in basic science coming from stem cell research and the sequencing of human genome as a result of the Human Genome Project completed in 2003 haven't resulted in clinical applications as soon as they were expected to (Solomon 2015, 161-163). The "pipeline

⁴² See Biddle (2007) for a discussion of the Vioxx case.

problem” is the term used to depict the slowdown, instead of the expected acceleration, in innovative medical therapies reaching patients (FDA 2004). What has been sought for is the “uncorking of the bottleneck” of pharmaceutical innovation. Furthermore, it has been estimated that it takes 17 years on average for research to reach clinical practice, which has been considered too slow (Morris et al. 2011). A demand for better results motivated the US National Institutes of Health (NIH) to issue the *NIH Roadmap for Medical Research* in 2004 to transform the way biomedical research is conducted, with translational research becoming a “buzzword” (Fishburn 2013, 487), a “mantra” (Maienschein et al. 2008, 43), “in vogue” (Fang and Casadevall 2010, 563), and moreover, “an imperative” (Harrington and Hauskeller 2014).

The Roadmap includes 28 initiatives in three main categories that broadly cover discovery, research, and implementation. The programs are “goal-driven, so that specific, high-impact outcomes could be reached within a set schedule of 5 to 10 years” (NIH 2014, 3). The cluster of initiatives focused on the implementation of new knowledge in the clinics is called “Reengineering the clinical research enterprise”, and is “central to the goal of moving research results more quickly into clinical settings” (Kantor 2008, 13). It includes Clinical research networks, Clinical outcomes assessment, Clinical research training, Clinical research policy analysis and coordination, and finally, Translational research (NIH 2014; Kantor 2008). NIH initially invested approximately 128 million dollars in the Roadmap program, with a total of 2.2 billion planned until 2009 (Williams 2005, 173).

A similar development in biomedical research politics occurred in the United Kingdom around the same time. The UK based *Journal of Translational Medicine* was launched in 2003 and *A review of UK health research funding*, known as The Cooksey Report, was published in 2006 (Aronson forthcoming). It discusses initiatives in line with the Roadmap and partly inspired by it. In the same year NIH launched the Clinical and Translational Science Awards consortium comprised of nearly 50 centers throughout the USA, with an investment of 500 million dollars annually by 2012 (Fang and Casadevall 2010). In 2009, the American Association for the Advancement of Science started the journal *Science Translational Medicine*. Soon translational research centers started emerging all over the world.⁴³

⁴³ For example, the 2016-2020 Strategy of the Clinical Hospital Centre Rijeka (KBC 2016) in Rijeka, Croatia (128.624 citizens in 2011) states several important and connected goals, among them: building new hospital buildings, founding a university hospital, and improving cooperation with different stakeholders, primarily the Faculty of Medicine in Rijeka. Especially emphasised is the 2008 initiative

The translational approach is based on the prospects of directly matching ideas for new therapies with the needs of patients observed in the clinic. The “bench to bedside” model has soon been updated to a bi-lateral two-way iteration process between discovery and the clinic (“bench to bedside and back”). Besides developing partnerships between laboratory and clinical researchers, clinicians, developers of medical care systems, and patients, the idea was to “fully involve and empower the public in the research process” (Maienschein et al. 2008, 44). The European Society for Translational Medicine (EUSTM) therefore adds community as the third pillar in the translational medicine model, along with bench and bedside (Cohrs et al. 2015). Miriam Solomon points out another feature pertinent to translational research, and that is physical proximity of preclinical and clinical laboratories, stating that it is “an interesting retro-invention in days of global electronic communication and global travel” (2011, 463).

The efforts around Roadmap initiatives, especially translation, present the most explicit attempt at speeding up research in order to achieve scientifically and socially important goals. Insistence on translation admits the difficulty of actually achieving it and signals the dimensions of crisis that had motivated the efforts to form a new systemized approach (Maienschein et al. 2008, 44; van der Scheer et al. 2017, 226). While many scholars have addressed the epistemology (Wehling 2006, 2008; Dougherty 2009; Fang and Casadevall 2010; Solomon 2015; Aronson 2017a, b, c, d, e; Robinson 2019) and ethics of translational medicine (Kagarise and Sheldon 2000; Maienschein et al. 2008; Sofaer and Eyal 2010; Kimmelman and London 2011), proposed suggestions for improved realization of its goals (van der Scheer et al. 2017), quantified its impact in current medical science (Zhang, Diao, and Wang 2014), engaged in field work with scientists who actually do it (Harrington 2011; Harrington and Hauskeller 2014), and critically evaluated different meanings of translational label (van der Laan and Boenink 2015), I focus on the possibility and practice of speeding up the process of discovery and research. In doing

started at the Medical Faculty of Rijeka to develop translational medical research, culminating in the successful project TransMed funded by the European Commission. The project is deemed exceptionally important as a way to gather preclinical and clinical researchers in cooperative work, which is done at the Centre for Translational Medical Research based at the University Campus in Trsat. This new location is hoped to be the centre of new research findings and clinical applications once the new clinic is built in the vicinity. In this way medical knowledge from the university can be directly applied to clinical practice in the hospital. What is most important, the two are planned to be brought together under the same institutional framework, at this point disconnected administratively, financially, and scientifically.

that I start from the assumption that acceleration involves coordination of research efforts towards particular goals motivated by possibly conflicting values (such as patients' safety vs. economic interest, regulation vs. early access, understanding vs. intervening), and developing new methods to achieve these goals. This part of the thesis aims to identify epistemic and non-epistemic means of and problems for an expedited translation of a treatment. Different means of acceleration will be introduced: some coming from external measures, while some being internal to scientific practice itself. Those internal to scientific practice involve trade-offs which will be spelled out and taken up for discussion in Part III of the thesis. The distinction between internal and external measures, however, will often be hard to maintain.

In the second chapter, the idea and different models of translational medicine will be presented, and the proposed ways in which translational medicine should contribute to the goal of accelerating scientific discovery and research will be introduced. In the third chapter, a historical overview of the reasons for the emergence of translational medicine will be given, in order to show how the goal to accelerate research came to focus. In the fourth chapter, some tensions inherent to the goals of translational medicine will be spelled out. In the fifth chapter, the analysis will be broadened to an examination of translational medicine in relation to a broader landscape of biomedical movements. This will be necessary in order to cover different stages of medical research and different epistemic trade-offs involved, and also to appropriately contextualize translational efforts. After that, in Chapter 6, some limitations to the goal of acceleration will be examined. The ethical considerations related to acceleration will be assessed in Chapter 7. Chapter 8 will then analyze an example of a paradigmatic case of translational research in the pre-translational era, namely research on cortisone, in order to illustrate some common limitations of translational efforts, but also the differences in the broader socio-economic context in which biomedical research was done several decades ago, as opposed to today. Chapter 9 will be the concluding chapter in which the broader socio-economic context of translational medicine will be highlighted again, both to conclude the translational case study and to set the stage for a discussion of values in the third part of the thesis.

2. What is translational medicine?

According to Jeffrey Aronson, by 1994 authors were starting to use the term "translational research" in the titles of their papers, then "translational medicine" in 1996, and "translational

science” in 2003 (Aronson 2017a). Miriam Solomon dates the emergence of translational research in the late 1990s, and translational medicine in the early 2000s (2015, 157). Mark Robinson also emphasizes that translational medicine is one of the many “monikers” (he adds clinical translation and molecular medicine alongside translational research/science) used to describe the larger aim “to turn basic science research done in laboratories into clinical applications” (2019, 3). I use the term “translational medicine” to keep the focus on biomedical science which this approach is supposed to transform.

There are many definitions of translational research, medicine, and science. A thorough discussion of these definitions can be found in “Jeff Aronson's words”, in British Medical Journal blogs. The definition I use in the introduction is a shorter version of Aronson’s definition of translational research in the clinical context:

“Translational research is the interactive use of the whole spectrum of scientific research, from basic to applied, to obtain new knowledge and understanding, and to develop new skills and innovative products or processes, any or all of which can be disseminated, tested, and implemented as monitored therapeutic interventions.”
(Aronson 2017e)

Miriam Solomon also discusses translational medicine definitions (2015, 157-159) but focuses more on metaphors like “bench to bedside“, “crossing the valley of death” or “Petri dish to people” (ibid., p. 164-166). She criticizes definitions for being too inclusive, allowing for “some jumping on the translational medicine bandwagon” (ibid., p. 159), and concludes that the term “translation” in this context is more aspirational than explanatory:

“It is primarily intended to inspire confidence that applied research will eventually be successful.” (ibid., p. 166)

Solomon identifies further buzzwords connected to translational medicine: “synergize,” “catalyze” and “interdisciplinary” (2011, 463). For example, an assessment of the Roadmap success, “NIH Roadmap/Common Fund at ten years”, concludes with the sentence:

“Although not every program has been completely successful, the *synergies* created by the Common Fund are inspiring, and the *catalytic* nature of these programs ensures that their impact will continue to grow.” (Collins, Wilder, and Zerhouni 2014, 276)⁴⁴

Most of the philosophical work on translational medicine shares the view that it is hard to “find substance amidst the rhetoric” and that the movement “appears to offer no more than a metaphor” (Fuller 2016). In order to identify and evaluate the meaning of translational label, Anna Laura van der Laan and Marianne Boenink (2015) have done a useful analysis of biomedical scientific literature on translational research/medicine. They have identified three dimensions along which translational label is usually discussed. The first dimension aims to identify the location and scope of translational gap(s), the second dimension models translational process, and the third dimension deals with the causes of perceived gap.

2.1 The location and scope of translational gap(s)

There are two widely recognized translational gaps, T1 and T2'.⁴⁵ T1 refers to translating ideas from basic and clinical research to the development of new products and practices, while T2' refers to implementing new products and therapeutic practices into clinical practice. To bridge the first gap means to apply discoveries generated in laboratory research and preclinical studies to the development of studies in humans in early trials (phase I and II). In phase I, therapies are evaluated for safety in healthy persons, while phase II trials seek for early signs of efficacy in a small sample of diseased patients. Van der Laan and Boenink's (2015) analysis shows that the focus on T1 is prevalent in most of the literature on translational medicine and they call it the narrow conception of the translational gap.

Bridging the second gap (T2') means enhancing the adoption of best practices in the community. A broad conception of the translational gap includes both T1 and T2'. T2' is represented in various ways, with fewer or more translations. In Aronson (2017c) a detailed identification of translational gaps is taken from the Tufts Clinical and Translational Science Institute (CTSI: “What is translational science?”, n.d.), according to which the translational

⁴⁴ Emphasis in the italics are mine. Elias Zerhouni, one of the authors of the review, was the director of the NIH in 2004 when the Roadmap was issued.

⁴⁵ T2' is the implementation gap understood broadly, while T2 is a particular gap inside of T2' between phase II and phase III clinical trials. The distinction will be explained further in the text.

stage that focuses on practice and implementation is comprised of four sub-translations (T2-T5) which respond to four perceived gaps.

First, there is T2 translation which expands discovery to larger patient populations in clinics during phase III trials. In phase III trial a novel therapy is evaluated against a standard therapy of choice (preferably) or against a placebo (if the standard therapy does not exist), and is supposed to show efficacy on a larger scale (sometimes also called “effectiveness”, though effectiveness is usually assessed in real world, non-experimental contexts). Phase III trials are usually randomized and double or triple blinded (masked) in order to reduce the effect of confounding factors (by randomization), and to reduce voluntary and involuntary bias in assessing the outcomes of a trial (by blinding). Randomization means that patients are randomly assigned to either the experimental or the control group. When a confounding variable (for example age, gender, general state of health or physical activity) is equally distributed among the two groups, then it should not affect the outcome of a trial. Blinding means that not everyone involved in a trial knows how patients are distributed in the two groups. In single-blinded studies the patient is not told what she is being given (a new therapy or a control therapy/placebo), i.e. in which group she is, while the doctors, medical staff and the pharmacist who assesses the results know about the dosage and the kind of substance that is being administered (a new treatment or a control). In a double-blind study, both the patient and the medical staff do not know what is being given (only the pharmacist knows), while in completely blind or triple masked studies, neither patients, nor medical staff, nor the pharmacist know what is being administered to whom. Biases can be connected to the selection of patients, dosage, randomization, and the selection of outcome measures and endpoints. Especially prominent is a tendency towards confirming the hypothesis that is being tested, i.e. the confirmation bias.

After phase III trials in T2, there is T3 translation that focuses on dissemination and implementation of therapies through the development of guidelines and recommendations, in order to learn how new interventions work in real-world settings. T4 then translates effective therapies and practices into improved community and population health care through the development of new policies. Finally, T5 adds a socio-political and economic dimension, involving public health education, prevention, and interventions into the social determinants of health.

An additional translational gap on the T1 side is T0 which deals with the basic science of technology development, including targets and biomarkers for developing interventions

(Aronson 2017c). Van der Laan and Boenink emphasize biomarkers as frequently suggested surrogate endpoints in translational research literature, developed to save time and costs of clinical trials (2015, 42-43). Examples of biomarkers are pulse and blood pressure, as well as complex laboratory blood tests and genetic tests. According to the NIH Biomarkers Definitions Working Group in Strimbu and Tavel (2010), a biomarker stands for “biological marker” and means “a characteristic that is objectively measured and evaluated as an indicator of normal biological processes, pathogenic processes, or pharmacologic responses to a therapeutic intervention” (463-464).

Surrogate endpoints, however, are often criticized as unreliable predictors of efficacy, especially in phase III trials (Strimbu and Tavel 2010; Stegenga 2015a, b; Holman 2019). For example, Holman (2019, 4366-4369) discusses a case of antiarrhythmic drugs that have caused a number of deaths in the early 1980s. The accepted surrogate endpoint for evaluating these drugs had been ventricular extra beat (VEB) which occurs when the left ventricle contracts before it has time to fill completely, so the heart fails to pump blood sufficiently. The antiarrhythmic drugs were shown to suppress the VEBs, but the mortality of the patients did not decrease. On the contrary, it had increased, but this was overlooked since the surrogate endpoint for assessing efficacy had been the suppression of VEBs, and not the more meaningful clinical endpoints like survival and death. Translational medicine makes use of biomarkers as surrogate endpoints because they are easy to measure and they can be measured relatively soon, but they nonetheless often fail in predicting the effectiveness of interventions.

Some literature on translational gaps conflates T2-T5 to one single broad stage of implementation and dissemination (T2’), but this is usually found in the earlier work on translational gaps, for example in the Cooksey Report. In more recent scholarly work, as well as on the webpages of translational centers and training programs, the particularization of two major gaps to further sub-gaps, especially in T2’, is widely recognized.

2.2 Three models of translational process

It is evident that the process of translation is primarily modeled linearly (basic to applied research), regardless of how the gaps are conceived. Translational efforts have mostly been informed by the linear model of science, though this model has often been recognized as

empirically inadequate due to the complexity of the innovation process.⁴⁶ According to the linear model, applications of basic knowledge were supposed to simply follow, naturally and unsupported, which was not happening as often as it was expected. The second dimension along which translational medicine is discussed is concerned with the models of the translational process.

According to van der Laan and Boenink (2015) there are three widespread models of the translational process: linear (bench to bedside direction), bi-directional (bench to bedside and back), and complex (interactive). Jane Maienschein et al. (2008) assert that the emphasis on translation suggests that if we want to use basic science with the aim of applying it, what we will be doing is translating or transferring from one context, institution, and scientific language to the other, which is recognized as an innovative process in itself.

“What is supposedly different here is an explicit recognition that translation is not easy, not inevitable, not unidirectional, and, indeed, not happening.” (Maienschein et al. 2008, 44)

Translational approach thus includes something more than the mere temporal sequence of distinct research practices, separated by goals, motivations, methods, institutional cultures, and disciplinary backgrounds of the researchers. The updated model “bench to bedside and back” captures the bi-directionality that is at work and presupposes “backward translation” which is able to channel the needs and values of patients to the basic research in the laboratory (van der Laan and Boenink 2015, 40-41). Bi-directionality also implies that the results from clinical trials are fed back into the design of new laboratory research (van der Scheer et al. 2017, 5).

The third, interactionist model, emphasizes the complexity of the innovation process and the feedback loops between different contexts, as well as across translational stages. Aronson’s definition explicitly mentions the “interactive use of the whole spectrum of scientific research”. In the traditional interactionist model the exchange between basic and applied is happening without the introduction of a specialized translational step. According to interactionism, the knowledge used in applied research has originated in overarching theories produced in the basic context, but is empirically adapted for particular, local purposes. Also, practically relevant knowledge can lead to fundamental issues which provide new understanding (Adam, Carrier, and Wilholt 2006). Contemporary translational medicine acknowledges that and is supposed to

⁴⁶ See Part I, p. 17-20 for a discussion of linear and interactionist models of research.

actively enable the interaction by multiple iterative steps. According to the complex or interactive model, translation is “a continuous data exchange within and between various research and non-research practices” (van der Laan and Boenink 2015, 40).

Much of the early criticism of translational medicine is grounded in its transitive status between basic and applied research, which are usually characterized by different goals and motivations. Most of the critical views grant more import for translation on the curiosity-driven basic side than on the application-driven clinical side, representing translational medicine as fundamentally misconceived. Ferric Fang and Arturo Casadevall (2010) for example, reclaim the linear model of scientific innovation against what they see as the imperative of application:

“It will be critical not to allow our impatience for translational applications to skew resources and researchers away from the open-ended exploration of the natural world that has provided the foundation for so many translational successes and remains as essential as ever.” (Fang and Casadevall 2010, 564)

Their concern is that the need for applications will obscure the value of curiosity driven research which they claim was the engine of therapeutic innovation in the past. They invoke the linear model of science as the role model and assert that insisting only on research aimed at solving disease-specific problems is a mistake, since “predictability is not the mantle under which we fund this stuff” (ibid.). They argue that it is impossible to design research towards predicted outcomes, since it is not known when genuine innovations will occur. Also, that scientific work is not driven by usefulness, but by curiosity, and that innovation happens in an uncategorized way. Modeling of the translational process is here closely related to the goals and motivations of translational medicine. Goals and motivations are one dimension not mentioned by van der Laan and Boenink (2015) because it relates more to the programmatic documents, mostly the NIH Roadmap, while their analysis focuses on biomedical research literature which engages with translational practice.

The line of criticism advanced by Fang and Casadevall (2010) is based on the normative claims of translational medicine and focuses on the appropriate means of achieving applications. We can thus call it the instrumental critique. According to the instrumental critique, the goal of application is thought to be better served by a strong base of open-ended research. The authors of *The Future of Drug Discovery: Who Decides Which Diseases to Treat?* are also in favor of this:

“The answer is not in having government funds directed more towards ‘translational research’. Basic research has to provide the targets for drug development.” (Bartfai and Lees 2013, 106)

Bartfai and Lees, like Fang and Casadevall are primarily concerned about the downplaying of the basic side in the translational process. Maienschein et al. (2008) (as well as Harrington and Hauskeller 2014 and Jogalekar 2012) share similar concerns, but they formulate them as a concern about what they deem is a dubious translational ethos. According to them, the problem is not only in the appropriate means of achieving applications, it is in the negative influence on the epistemological practices that comes from overemphasizing the application goal. The imperative of translation towards application, they believe, might be altering the research process itself.

“Where the Vannevar Bush model emphasized purity of science in its own right (even when it could also have worthy applications), today’s translational research builds certain (and sometimes dubious) end goals into the research from the start. One lesson is that assuming outcomes (however well-intentioned) alters the research endeavor.” (Maienschein et al. 2008, 49)

Their criticism is also a version of the instrumental critique, but spelled out from both an epistemological and an ethical perspective. What Fang and Casadevall are saying is: yes, we want applications, but this is not how we get them. What Maienschein et al. are saying is: we do not want applications at any cost, and in the case of translational medicine, it is to be expected that the cost is both epistemological and ethical. The epistemological critique is that science is distorted by a predominant focus on application at the expense of understanding. The ethical critique is that the acceleration of research and development comes at the expense of a stalled or avoided ethical discussion.

“The problem is taking the translation as an unquestioned desirable goal and trying to make the ethics fit. This distorts the ethical discussion as well as the science.” (Maienschein et al. 2008, 50)

For example, an exclusive focus on “what works” might be risking the occurrence of biases in the research design that favor false positives, which might lead to an overestimation of effectiveness and an underestimation of harm. A further problem is pursuing research on applications and therapeutic interventions without adequate knowledge of the underlying mechanisms. Maienschein et al. (2008) put forward this argument as the main reason for caution

about stem cell research in particular, which they consider “a poster child for translational research” (2008, 48). In stem cell research, they argue, we are “translating from sadly incomplete benchside and bedside source languages, languages with unknown grammar, unknown syntax, and few if any native speakers” (ibid.). Ashutosh Jogalekar (2012) gives a similar, more general conclusion: “Committing national resources and public attention to translational research when most of the basics are still to be understood is an endeavor fraught with great risk and uncertainty.”

Jean Harrington and Christine Hauskeller (2014) have carried out an ethnographic study on stem cell research from 2006 and 2011 with the aim to explore the motivations and attitudes of the researchers working in this field, and found out that conforming to the “translational imperative” was necessary in order to get funded, even when one did not know if applications are likely to occur:

“So I have to play the game, I have to play the rules of the game because in the end what I want is to be funded and to be in a lab working and doing research. (...) So... it’s more about, [pause] giving the people what they want to read, even if inside you know it’s not necessarily achievable, or it’s not your first priority, but again you have to combine all these things, basic research with translational research and get the money.” (A report of a stem cell researcher in Harrington and Hauskeller 2014, 79)

According to this report, translation is more of a marketing label and a declarative prerequisite for getting funding. This dynamics definitely poses ethical concerns since it involves a deceit, but it at least reassures that basic work has to be done no matter under which name. Finally, the most problematic aspect of translational ethos is that the well-intended goals might be instrumentalized as a reason to justify getting around the ethical discussion of the emerging technologies, and rush to applications.

Harrington and Hauskeller draw a very similar conclusion:

“That research has to be oriented toward therapeutic application to deserve public funding and be of societal value is an imperative that contradicts and challenges to the point of denial the complexity of successful interactions and transfers between multiple agencies. Biomedicine is pregnant with translation.” (Harrington and Hauskeller 2014, 80)

What they are saying is: science has been translational before and there is no need to change how innovations have been achieved. Overemphasizing goals can only compromise them. I will stay on the epistemological side for now, and turn to the ethical problems of translational medicine later. I would like to point out an interesting inversion of positions when it comes to attempts at reclaiming the linear model, especially in Fang and Casadevall (2010).

Fang and Casadevall see translational medicine as a formalization of the application-driven interactive model of science and set out to defend the curiosity-driven linear model against it, believing that it is flawed to base science funding on erroneous understanding of the innovation process. They understand the application-driven interaction of basic and applied as the new norm, and an empirically inadequate one, since there are many examples of clearly linear translations.⁴⁷ It is arguable, however, whether interaction is really as opposed to linearity as it may seem. It depends on whether the models are understood descriptively or normatively. Unlike the linear model which was established as a norm in 1945, interactionism emerged as a descriptive model (see Adam, Carrier, and Wilholt 2006) and its merits were recognized in translational initiatives. Now interactionism occurs as a norm in translational medicine, while the linear model is evoked to describe how translations are actually achieved. The situation is inverted, and Fang and Casadevall raise flags against the normativity of the application-driven interactionism established by the translational paradigm.

But interactionism is not necessarily against linearity, it only states that linearity often does not work, and that it might need to be supported. It is against linearity as a norm, but it can descriptively be successful. Also, the linear sequence of events (from discovery to application) can hold even when the institutional context is more diverse and interactionist (from clinical research to laboratory research, i.e. “backward translational”). Motivations are even harder to tie strictly to one kind of outcomes (curiosity to understanding, problem-solving to innovation). Allowing for more diverse institutional contexts and iterative steps does not deny any direction of translation, so criticism along the lines of the linear model as empirically adequate is misguided. The linear model can be seen as a subset of the interactionist model, a possibility of interaction executed linearly in a particular case. The biggest problem with the linear model as

⁴⁷ The first example they give is Marie Curie’s discovery of radium which was not application-driven though radioisotopes soon gained wide therapeutic use in hospitals. Another example is the research on insect embryogenesis that has led to advances in innate immunity and development of drugs and vaccines (Fang and Casadevall 2010, 565).

presented in 1945 is its normativity, the idea that this is how innovation should work, similar in strength to what Harrington and Hauskeller (2014) have called the “translational imperative” and Maienchein et al. (2008) the “translational ethos” in the new context. When uttered as descriptive models, both do justice to certain cases of innovation, but do not cover all of them.

Contrary to these critical views, the basic side of translational process is sometimes recognized as the strength of translational models. Miriam Solomon (2011, 2015), for example, understands translational initiatives as primarily focused on understanding disease mechanisms that are central to basic research. A focus on application, especially through an emphasis on randomized controlled clinical trials, is according to her a characteristic of the evidence-based paradigm in medicine. In fact, she identifies the limitations of evidence-based medicine as one of the reasons for the rise of translational initiatives which reclaim the quest for understanding as opposed to reliance on statistical methods. This is quite the opposite view than the view of Fang and Casadevall and others who see translation as primarily about application, sometimes at the expense of a lack of understanding.

Finally, according to the third view, translational medicine is just ‘business as usual’ without any particular drive for application apart from a declarative and aspiratory one, as we could see from the stem cell researcher’s report. This is by no means a stand-alone view. It has been suggested, for example, as early as 2006 in a paper with the title “Translational science – a sexy title for pre-clinical and clinical pharmacology” (Johnstone 2006), and more recently it is a view shared by Mark Robinson (2019). Solomon (2015) also emphasizes the aspiratory role of the label.

Translational medicine’s declarative role is closely related to the next issue that focuses on identifying the causes of translational gaps and means of addressing them. The question is: Is translational medicine really a new kind of medical science? Is there anything internally, epistemically different about it? While the short answer will ultimately be “no”, there is still a more complex story to tell about why this is so, and what is really new on the horizon of biomedical research practices.

2.3 Causes of translational gap(s)

The third dimension along which translational medicine is discussed aims to identify the causes of perceived translational gap(s). Van der Laan and Boenink (2015) identify two groups of

causes: external and internal. External causes are non-scientific and come from the context in which science is practiced, while internal causes have to do with the biomedical research practice itself. The literature on translation mostly centers on the external causes of translational gaps and does not aim to offer a genuinely new practice. It is rather aimed at providing external support to the existing ways of doing medicine and at removing obstacles that stand in the way of efficient translations. These obstacles can be practical, social, economic or ethical. For example, an obstacle can be too much administrative work (practical), lack of communication between laboratory researchers and clinicians (social), lack of funds for expensive clinical trials (economic), or strict regulations for research with human subjects (ethical) (van der Laan and Boenink 2015, 41-42). The proposed measures can thus be: removing regulatory obstacles, changing the recruitment of research subjects, finding new ways of training researchers, publishing preliminary research data, developing clinical guidelines, or stimulating collaboration between academy and industry, as well as between researchers and clinicians (ibid.).

Contrary to that, a view of translation perceived as an internal, scientific problem focuses on scientific methods used in scientific practice. Problems occur at different stages of translation. For example, it has been argued that *in vitro* and animal models produce a type of knowledge that does not adequately represent the complex mechanisms of the human body, so it often fails to translate in T1. What is needed are more complex animal models or models based on human tissue. On the clinical assessment level, the gap is thought to be increased by the experimental set up of randomized control trials, which do not translate well to the real-world settings and are often successful only for a small group of patients (ibid., p. 42-43).

“In documents identifying ‘internal’, scientific causes for the lack of translation, translational research is anything but ‘science as usual’. (...) In addition to producing more realistic research models and using real world data, a recurrent theme is the need for convergence or (re-)integration: of different life sciences, of different experimental approaches, of life sciences and clinical sciences, and even of life sciences and all kinds of population studies. Such integration requires computational research with large databases of molecular, clinical and epidemiological information. Better information technology systems are therefore assumed to be a critical condition for translational research – indicating that these novel research methods, like the old ones, should be facilitated.” (ibid., p. 43)

Better integration of various data sets and the usage of computational methods, among else, have eventually led to an influential and genuinely internal scientific transition, and that is a transition towards personalized medicine. Since this part of the thesis aims to identify epistemic means and ethical problems for an expedited translation of a treatment as a part of the overall aim to explore the ways in which scientific research deals with different time constraints and their relation to goal setting, design of methods, and evaluation and communication of results, it will be necessary to engage with a broader set of medical movements and practices in order to address these relations. Although translational medicine is very explicit about acceleration as its goal, the predominantly external measures of acceleration which it offers will not be enough for a meaningful epistemological and ethical analysis. The “personalization” of medicine in recent years presents the more genuine internal change in contemporary biomedicine and it is important to see how translational medicine has facilitated the transition towards it. Furthermore, it will be shown that it is very hard to delineate the internal from the external causes of translational gaps, as well as to distinguish the respective measures to address them, so the distinction itself will often be hard to maintain. Van der Laan and Boenink also admit that the border between science and its context is “porous”, and that the “domains of external and internal causes of translational gaps often overlap” (2015, 43).

In the next chapter an overview of the history of translational medicine will be given. This will allow us to see which processes and political decisions have led to its emergence as a formalized approach, i.e. which factors have led to the prevalent focus on acceleration of medical research towards applications and implementable products. This will also make clear certain tensions and conflicts of values and interests of several stakeholders involved in medical research, which will be spelled out in Chapter 4. Chapter 5 will then be devoted to an exploration and evaluation of translational medicine in the context of a broader landscape of biomedical practices. Two other prominent approaches in biomedical science will be introduced: personalized medicine and evidence-based medicine. Only through an understanding of what the contemporary biomedical research is comprised of can we assess the role of translational medicine inside of it and comparatively identify its characteristics. More context is a prerequisite for evaluating whether the internal/external cause distinction holds, and whether and how have translational practices addressed it. It will be argued that the emergence of personalized medicine is closely related to translational efforts and that accelerated translations are nowadays most often achieved on the terrain of personalized medicine. Also, that certain features of evidence-based medicine, on the other hand, do not align well with translational efforts, but that they also do

not share the same focus. This will highlight the trade-offs inherent to medical practice in general.

Finally, the complex interactions taking place in contemporary biomedical landscape cannot be adequately described and assessed independently of social, political, and economic factors whose interests and values do not always overlap. On the contrary, they can be a source of significant conflicts which can influence the way research is conducted and evaluated. A necessity of prioritization among different aims and goals of research will become clearer, as well as trade-offs between different values pertaining to research practices, which will be taken up in the last part of the dissertation.

3. A short history of translational medicine⁴⁸

Medical science has been translational from its very beginning, but translations have been achieved in an informal and unsystematic way. In the Cooksey Report, translational research is exemplified by the case of discovery and clinical application of penicillin. Alexander Fleming's discovery in the laboratory could not have been used in patients without its successful translation by Howard Florey, Ernst Chain, and Norman Heatley. However, the laboratory work was in this case done in the hospital, while translational work was done at the university. Fleming worked at St. Mary's Hospital in London, while Florey and Chain worked at the Oxford University. The traditional linear model assumes the opposite – basic work is/should be done in the curiosity-driven context of the university. The translation was achieved linearly to a certain degree, since the direction was from a serendipitous discovery in the laboratory to its application in the clinic, but the basic research was done in the application-driven context of the hospital, which is in line with the interactivity suggested in the contemporary translational approach where clinical needs inform the laboratory work. It is also in line with my view that the linear sequence of events (from discovery to application) does not have to be opposed to

⁴⁸ I am indebted to Jeffrey Kenneth Aronson for providing me with invaluable insight into the history and the concept of translational medicine by sending me his draft on translational research from the forthcoming textbook on molecular medicine, part of which has been published in his blogs on the British Medical Journal webpage. Our conversation and subsequent e-mail exchange following a conference in Munich in April 2016 has been crucial for my involvement in the topic. In this chapter I am following the historical development as outlined in his draft “The translation of pharmacological actions of medications into clinical outcomes” (Chapter 2.8).

the interactivity of the context (from hospital to the university), also that either linear or the interactive view have to be related to any exclusive motivation to start from. Especially in the medical context, epistemic and social motivations are very hard to disentangle.

In 1928 Fleming observed mold which inhibited the growth of bacteria on plates contaminated with it. By 1931, however, his efforts to translate this discovery to treating infections in animals via oral administration resulted in failure, and he concluded that the substance could not stay in the body long enough to be effective. It was in 1939 that Florey, Chain, and Heatley managed to purify small quantities of active penicillin and successfully treat infected mice with it. Pressed by the wartime circumstances, they tried to scale-up penicillin production using bedpans, milk churns, and bathtubs as culture vessels. By 1941 they developed techniques that allowed production of sufficient quantities to start tests on humans and soon they transferred their research to the US Department of Agriculture laboratories, in order to find a way to produce penicillin on a larger scale. It was the Merck Company that supplied first larger quantities to treat the injured soldiers. In 1945 Florey, Chain and Fleming shared the Nobel Prize in physiology or medicine for their work on penicillin, and in 1948 Andrew J. Moyer was granted a patent for the method of mass production of penicillin. (Cooksey 2006, 17)

The story of penicillin is a case of a successful translation in the pre-translational era. The Second World War created conditions which had accelerated the testing of many treatments and the pace of progress was advancing at an unprecedented scale. The idea of translational research as a formal process, however, can be traced back to the idea of diffusion of innovations and technology transfer (Aronson forthcoming, 3-4). The US Stevenson–Wydler Technology Innovation Act of 1980 facilitated technology transfer from federal laboratories to nonfederal entities and it provided access to federal laboratory technologies to outside organizations (see Jolly 1980). In the same year, the Bayh-Dole Act enabled universities, nonprofit research institutions, and small businesses to own, patent, and commercialize innovations developed under government funded research programs within their organizations. The act allowed and incentivized more universities to become actively involved in the transfer of technology from the laboratory to the market, with the aim of advancing translation (see Loewenberg 2009). Under the Federal Technology Transfer Act of 1986 all laboratories are expected to improve transfer activities and to focus on firms that will commercialize technologies. Responsibility to promote technology transfer was also reflected in laboratory job descriptions, employee promotion policies, and job performance evaluation (Bagur and Guissinger 1987, 53).

Nonetheless, the efficiency of new drug development in the 1990's and 2000's has been suboptimal due to the separation between those bringing about the discoveries needed for new therapies, and those with the funding and commercial capabilities needed to bring the drugs to the market (Fishburn 2013, 487). The ideas for new therapies have usually arisen in academic institutions or biotech companies which are able to perform only early-stage research before needing to raise money from investors. Normally it is the large pharmaceutical companies that would then take on and finance randomized phase III clinical trials, file submissions to regulators, and perform the sales and marketing (ibid.). The flow of knowledge-product development from the academic to the industrial culture, and further into clinics and patients' hands, was very insecure and often not happening even for promising therapies.

Reasons to insist on translation have been social, political, and economic. Fang and Casadevall have stated that translation is an easy sell, but which also fills a genuine need (2010, 563). This need comes in part from "perverse incentives that value basic science more highly than applied research" (Cooksey 2006, 1). We can relate this qualification to the already mentioned recognition of the limitations of the linear model. Since applications were not following the advances in basic science, translation had to be incentivized and re-invented. The Cooksey Report states that translation had already been a norm in the UK in 2006 and that its predecessor had been technology transfer, in the UK context as well as in the US (2006, 20).

In the year of the Roadmap, FDA also issued *The Critical Path Initiative*, national strategy for transforming the way medical products are developed and evaluated, as well as a publication called *Innovation/Stagnation: Challenge and Opportunity on the Critical Path to New Medical Products* (FDA 2004).

"Sounding the alarm on the increasing difficulty and unpredictability of medical product development, the report concluded that collective action was needed to modernize scientific and technical tools as well as harness information technology to evaluate and predict the safety, effectiveness, and manufacturability of medical products." (FDA 2018d)

The development of new medical products, especially pharmaceuticals, was considered to be too long, too expensive, and altogether too inefficient. The Critical Path involves prototype design or discovery, preclinical development, and finally clinical development (stage I-III clinical trials), at the end of which a company files for market authorization, drug approval and launch preparation (Cooksey 2006, 106). Further stages are those of assessment of clinical- and

cost-effectiveness, and stage IV clinical trials which can either be addressing particular issues raised by a regulator or assessing how new products are used in practice (ibid.). The last stage is the ongoing safety monitoring. The process is characterized as a very linear one, but with uncertain end-points (ibid.). The following chapter summarizes the problem from a policy maker's point of view:

“Clearly, this is an extremely complex and long process. One estimate is that a new drug typically takes 12 years to reach the stage of being given a marketing authorisation (although device development is typically shorter and more incremental).⁴⁹ Given that patent life lasts 20 years, this clearly raises challenges for pharmaceutical companies in terms of recovering the cost of developing the new drugs and making a profit. This perhaps makes it unsurprising that many companies are worried about developments such as the increasing use of technology assessment processes, not only because it makes judgements about clinical- and cost-effectiveness that could reduce access to markets, but also because of the time involved in reaching even a positive judgement (this is just as true for manufacturers of medical devices, even though patents are not always as crucial an issue as they are for pharmaceuticals). This can also be frustrating for patients who, understandably, want access to new treatments as fast as possible, particularly when existing ones are of limited effectiveness.” (Cooksey 2006, 106)

Several issues have been raised: the complexity of healthcare product development and its regulation, the evaluation arguably coming late in the process, the constraints posed by patent law, and inability to gain profit. The complexities increase the time it takes to develop a new drug or device, therefore increasing the cost directly by higher costs, and indirectly by shorter actual patent life. However, these are primarily producer-oriented issues. Patients are introduced as a party that surely would not mind to see some changes, while the major source of worry seems to be the regulation and increased costs of research and development with no return for the producers. A note is needed on the current state of pharma industry and its consolidation in the face of crisis.

⁴⁹ The Reports adds a note that the 12 years average is taken from DiMasi et al. (2003, 164, 181), and it depicts time taken to develop a new molecular entity, which is a category that 35 percent of new drugs fall into. The issue of translational lag measurement will be discussed further in the text, especially since Morris et al. (2011) meta-analysis shows that several studies converge around 17 years as the average time of translating research evidence into therapy.

The 1970's were the start of the blockbuster era, with the first blockbuster drug, cimetidine, launched by GlaxoSmithKline in 1977, after it was discovered in 1971. Cimetidine, sold under the brand name Tagamet, inhibits stomach acid production and is used in acid indigestion and peptic ulcer. Ten years after its introduction, it had achieved sales of one billion dollars and became the world's leading prescription drug, thus becoming a blockbuster – a drug that generates annual sales of at least 1 billion dollars. Cimetidine also represented a revolution in the way pharmaceuticals are developed, since it was one of the first drugs to be designed logically, by “rational” or “structure based” drug design (Taylor 2016, 23; “American Chemical Society: Tagamet”, n.d.), as opposed to a serendipitous discovery.⁵⁰

For the next decade, the blockbuster model proved to be successful, with other pharmaceuticals being regularly patented, and the profits made during their patent life of twenty years sufficient to fund research and development of future products (Taylor 2016, 23). Blockbusters are sometimes referred to as the “low-hanging fruit” of pharmaceutical innovation, since they were relatively easy to discover and turn into marketable drugs used widely in population. They were also more easily approved because they were usually tested against a placebo, which was the standard of care when they were introduced (Bartfai and Lees 2013, 176). Today efficacy is often assessed against the existing treatment of choice, which makes it harder to establish a significant improvement. For example, all diabetes drugs have to work on top of metformin, and hypertension drugs on top of propranolol, which already have a well-established efficacy and safety record (*ibid.*). The standard of care is simply better than it was in the 1970s so it is harder to find new molecular pathways that would work on such a large scale (see Stegenga 2018). Once the blockbuster model was exhausted, a decline came, following a long-lasting productivity crisis in the pharma industry. However, that does not mean that companies have not taken measures or were not making money anymore.

The initial response to the problems was consolidation, with a number of smaller and bigger mergers and acquisitions of companies. By 2010 the 30 research pharmaceutical companies that existed in 1989 had merged to become only 9 companies (Taylor 2016, 24). The industry has adopted several strategies, such as significantly reducing operating costs, but also focusing on biotechnological innovation through the acquisition of biotech companies (*ibid.*, p. 23-25). In the 1980s pharma industry had started with the development of generics and me-too drugs, and more recently it has turned to biopharmaceuticals (biologics) and biosimilars, i.e. biological

⁵⁰ For an overview of the emergence of rational drug design in the pharma industry see Adam (2005).

therapeutics emerging from advances in omics research and systems biology, and related to the emergence of personalized medicine (see Trøst Jørgensen 2008). Omics technologies are aimed at the detection of genes (genomics), mRNA (transcriptomics), proteins (proteomics) and metabolites (metabolomics) in a specific biological sample in a non-targeted manner.⁵¹ The integration of omics technologies is characteristic of high-dimensional biology or systems biology because it relies on a holistic view of the molecules that make up a cell, tissue or organism (Horgan and Kenny 2011, 190).⁵² The omics field is enabled by technological advances in high-throughput analysis of biologic molecules which make it possible to quantify the levels of protein coding transcripts in a particular tissue, map loci that control gene expression, as well as to analyze a large amount of transcripts and metabolites simultaneously (Hasin, Seldin, and Lusic 2017, 1).⁵³

The Hatch-Waxman Act of 1984, officially called the Drug Price Competition and Patent Term Restoration Act allowed drug companies to produce generics – drugs with the same active ingredient and the same therapeutic effect as the brand name drug, but marketed after the original drug's patent has expired. Generics are usually less expensive than the brand name drugs because clinical trials for effectiveness and safety do not have to be duplicated (Sokal and Gerstenblith 2010; Mason 2013). A “me-too” drug or follow-on is similar to an existing drug, employing the same mechanism but with a different molecule and therefore a different patent (Jena et al. 2009). Biologics are medications produced from biological sources such as blood, proteins, sugars, or tissues, i.e. originating in living organisms. A biosimilar is a biological product that is very similar to a reference biological and for which there are no clinically meaningful differences, hence it is a version of a me-too medication (Mason 2013, 7). Clinical evidence for biosimilars, however, has to be assessed individually, as opposed to generics which share clinical data with the original product (ibid). The first biosimilar, Zarxio, was introduced on the market in 2015 (FDA 2018b). Biologics and biosimilars are produced through complex biotechnological processes, unlike chemical drugs and generics which are produced by chemical synthesis. During the period from 1998 to 2002 the FDA approved 415 new medications, out of which only 14% were new innovations, while 9% were significantly improved old drugs and 77% were as good as the existing ones, with no particular

⁵¹ I will say more on targeted and non-targeted methods in drug discovery in Chapter 6, p. 107-113.

⁵² I will say more on systems biology in Chapters 5 and 7.

⁵³ I will say more on high throughput screening in Chapter 7.

improvements (Brown 2017, 152). In any case, pharmaceutical companies have continued to make significant profit.

As a possible solution to the state of affairs identified earlier, namely the pipeline problem, and with an eye on multiple institutional, financial, and cultural obstacles for translating research into practice, the Cooksey Report suggests that:

“(…) the government, regulators and industry create a new partnership to pilot a new drug development ‘pathway’ to create wins for all stakeholders: industry, government, the wider economy and, most importantly, patients. This pathway should enable:

- more *rapid* discrimination between potential new therapies at earlier stages of drug development;
- *earlier* ‘conditional licensing’ of new drugs;
- involving NICE *earlier* in the process of development to *accelerate* assessment of clinical and cost-effectiveness;
- *faster* uptake of cost-effective drugs;
- *clearer* processes for ensuring NICE initial assessments and recommendations for further research are followed-up more *systematically*;
- the use of the NHS National Programme for IT (NPFIT) to ensure more *rapid* assessment of any emerging side-effects and efficacy over longer periods;
- *streamlining* of processes involved in setting up and costing clinical trials; and
- the use of NPFIT to identify appropriate patients for clinical trials” (Cooksey 2006, 6)⁵⁴

The new pathway is supposed to bring wins for all stakeholders by being “rapid”, “earlier”, “faster”, “accelerated”, “clearer”, and “streamlined”. The call for collective action towards translation means reengineering medical enterprise on multiple levels: institutional, political, cultural, economic, and epistemological. Though it seems obvious that translation is both an

⁵⁴ NICE – National Institute for Health and Clinical Excellence, NHS – National Health Service, NPFIT – National Programme for Information Technology. Italics are mine.

easy sell and a fulfillment of a genuine need, what is harder to discern is: for whom, why, and how?

4. Accelerating discovery and research – hopes and tensions

Fang and Casadevall utter the “genuine need/easy sell” phrase to depict how the US NIH responded to political pressures to demonstrate public benefit from scientific research when confronted with multifaceted problems in the information flow between basic science and clinical medicine (2010, 563). However, the stakeholders in the case of biomedical research are diverse. First, there are patients, who are in genuine need by having health issues unaddressed, or addressed not soon enough, or lacking access to new treatments, or suffering from side-effects of the existing treatments. Second, there are producers, who want to cover the costs of research and development. Third, there are regulators, who need to strike a balance between overregulating and under-regulating. Fourth, there are researchers, who might not be motivated by only usefulness, but rather by curiosity. However, academic work is being increasingly sponsored by industry in the new partnership, so the intrusion of market forces is not only tied to industry. The question suggests itself: which side is portrayed as being in genuine need, and which one is buying an easy sell? Patients are undoubtedly in need for better clinical outcomes, while translational initiative is an easy sell to funders and producers since they will benefit from speeding up. Or is it, on the contrary, in the producer’s primary interest to recover the costs and regain the sunken productivity, while the idea is easily sold to the public, including patients, who want to improve their health rather sooner than later?

This seems to be the crucial ambiguity, and a well-recognized one. Jonathan Fuller (2016) and Naomi Scheman use “bench to market” metaphor instead of “bench to bedside” to state exactly this tension.⁵⁵ Still, the goals of multiple stakeholders do not have to necessarily be opposed. One of the reasons is that these groups intersect. For example, everybody is a patient, especially with the increasing age of the population, which is the biggest risk factor for developing a health condition that needs to be treated. Second, because of the new attention to prevention coming from emerging technologies that can screen individual genetic setup for particular health issues. However, a problematic aspect of certain preventive practices can be over-diagnosing for

⁵⁵ Scheman used it in her talk at the 2018 Philosophy of Science Association conference in Seattle.

conditions that do not actually need to be addressed, especially when risks posed by preventive treatments exceed their benefits (see Biddle 2016).

Nonetheless, we have reasons to believe that improved clinical outcomes are in everybody's interest: even the "evil pharma", as it is often depicted in the public discourse, can want safe medications and fewer law suits. One might object that this is a somewhat naïve statement, like saying that industrial managers can want a world without climate change and that this is why they should change their CO₂ polices. The industry, or pharma, when representing the sector, is not a sum of its staff, nor is it characterized by the particularities of individual leaders that happen to be in charge at a certain time. It is fundamentally profit driven. Still, health care has been improved in the course of the last century and one of the reasons for this is the industrial production of live-saving and quality-of-life improving treatments. It is the production of these treatments in particular – safe, effective and broadly used – that has been cost-effective, and it has been expected and wanted that it continues being so in the future. However, this is somewhat unlikely taken that the "low hanging fruit", i.e. medications working as "magic bullets", targeted specifically at the causes of diseases and effective for a broad range of people, have already been discovered (see Stegenga 2018), especially in the blockbuster era. As a response to that and in order to get their products approved and widely used (sold), pharmaceutical companies have been associated with problematic research and dissemination practices (see Brown 2004; Biddle 2007; and Wilholt 2009). These problematic practices led to safety issues with pharmaceutical products, eventually decreasing public trust and increasing doubt about the presupposition that improved health outcomes really are a shared goal (see de Melo-Martín and Intemann 2018, 96-115).

Have clinical outcomes been improved as a result of translational medicine? Not long after the Roadmap was launched, we could read that "translational research generated revenue, brought publicity, and enhanced public relations" (Fang and Casadevall 2010, 564). No mention of clinical outcomes though, and it is not a surprise for at least two reasons. First, translational success is a complex interaction of multiple factors supposed to enhance the research and development process, which is recognized by advancing through the milestones on the way to new products and ultimately introducing a certain product on the market (early enough). However, this might or might not correspond to an improvement in diagnosis, prevention, or treatment of a disease or a condition that this product or practice is supposed to address.

“The success of translational research is not only a function of the quality of the science, but also of the collaboration between academia and industry, the organization and management of research and development (R&D), the public policies that regulate scientific research and the connections among the key people involved.” (Sánchez-Serrano 2006, 107)

The collaborative success will be recognized relatively early, in 5 to 10 years in which programs should be streamlined towards their envisaged endpoints on the critical path of delivering medical products to the market. However, their overall success in treating ailments and diseases will need to stand the test of time which extends well beyond reaching phase II and III trials, or even regulatory approval.

This brings us to the second point when it comes to evaluating clinical outcomes of translational medicine. A comprehensive picture of the effectiveness of a new therapy and its ability to bring about an actual change to the better, simply has to be waited for, usually for decades. It has to be shown whether the patients’ genuine needs have been fulfilled alongside the improved approval rates of the manufactured therapeutics. Regulatory approval is not a sufficient indicator, which has been learned the hard way from cases like thalidomide and Vioxx. The recent case of Vioxx is especially alarming since it is a case of an approved drug that had to be removed from the market because of its devastating effects that have made it through the regulatory process due to faulty and misleading research design.

Furthermore, we are currently in the second decade of the translational hype and though many products have reached the market, it is left to be shown to what degree, if at all, this is the result of re-engineering of the way clinical research is done, or it would have occurred anyway due to recent advances in basic science and increased expenditures. The translational direction that has been taken includes multiple pathways to biomedical innovation and its implementation, from collecting databanks and conducting omics research to enhancing implementation practices through guidelines and outreach. Translational medicine addresses a variety of contemporary ailments, from the most widespread like cancer, Alzheimer’s and depression, to research into rare diseases and orphan drug development. From an epistemological side it makes sense to ask how the acceleration of the research process is supposed to happen and whether translational practices involve anything genuinely new, apart from reallocating resources and incentivizing exchange between basic and clinical researchers.

Concerns arise on the ethical side as well. A number of studies have shown a correlation between shorter research time and subsequent safety problems, for example in trials testing the effectiveness of methylphenidate (Ritalin) to treat attention deficit hyperactivity disorder (ADHD) in children (Schachter et al. 2001, Molina et al. 2009). The treatment has shown small positive effects in the short run and the trials did not last long enough to see the long term effects that proved to be non-beneficial and even harmful, causing decreased body height and mass (see also Stegenga 2015b, 4). Shorter FDA approval time has also been associated with safety problems (Carpenter, Zucker, and Avorn 2008), which in Carpenter, Zucker, and Avorn's study is explained by the effect of imposing deadlines on regulation in combination with sub-optimal number of staff. However, another source suggests that the reason for increased safety problems coupled with shorter FDA approval times is the industry involvement through the user fees enacted by the PDUFA in 1992 (Horton 2001). Horton's argument is backed up by a 1998 survey of FDA medical officers who have reported that the standards for drug approval have declined since the PDUFA enactment:

“Many officers felt under greater pressure from FDA supervisors to approve new drugs; they received inappropriate calls from the sponsor about the drug under review; and they believed that the FDA too often interfered on the company's behalf in the drug-approval process” (Horton 2001, 1545)⁵⁶

On the other hand, patient advocacy groups have entered and influenced health policy in the past three decades, demanding more government research funding, quicker drug approvals, and improved access to experimental treatments, which is very much in line with the road that NIH and the Cooksey Report have taken. For example, the first drug for AIDS, azidothymidine (AZT), was approved more quickly than subsequent therapies, in part because of the pressure for quick approvals coming from patients' advocacy groups, especially because no effective therapies were available (Epstein 1996). Control groups were also avoided so that more patients could get the medication immediately. Epstein's study is precisely about the democratization of biomedical research through the involvement of patients as activists and partners in research, which is prior to and in line with the contemporary “Science in and for the society” idea, as well as with the inclusion of community as the third pillar of translational research. However, AZT was not as successful as it was first thought. A three year follow up study of its effectiveness conducted on two thousand patients showed that patients in the placebo group

⁵⁶ Interestingly and unfortunately, this survey is not available online anymore.

were more likely to survive the three years of study than patients on AZT, and that the drug had serious side effects and almost no benefits after a certain period of usage (Crewe 2018). This showed that quick access and avoidance of the control groups during the research phase was a mistake. It was later shown that AZT has beneficial effects but only in combination with other medications.⁵⁷

We can see from the cases of Ritalin and AZT that the increased speed or shorter research/evaluation time is highly problematic regardless of the motivations behind the acceleration. Ritalin is not a life-saving drug and it is especially controversial because it targets children. The medicalization of ADHD is problematic in itself, since the harmful effects of the condition are largely dependent on the social context. It is thus more likely that the shorter trial time was motivated by decreased cost of research and easier and quicker access to the market in the interest of the industry, because negative effects will not yet manifest by the time of the evaluation. The case of AZT, on the other hand, shows that accelerated research practices and quicker approval can also be motivated by hopes of benefit for the critically affected population. Finally, acceleration can also be a consequence of imposing unrealistic deadlines without adequate resources, regardless of whether this is done in good or bad faith. In all of these cases, shorter research time is coupled with subsequent safety problems.

On the other hand, the Ebola case gives support for the exactly opposite claim. Accelerated trials and earlier access have managed to control the epidemic. Evidently, we want speed when it provides access to a better state of affairs and we do not want it when it brings about the worse. It seems to be a leap of faith on which side we end up, and this can only be ascertained in hindsight. Nonetheless, there are some systemic factors that need to be accounted for in the recent attempts at accelerating. The focus will not be on the speed of FDA regulation⁵⁸ or controversies about the patent law and the role of regulatory agencies in general (see Brown 2017; Reiss 2017; and Biddle 2013a). What will be analyzed in the following chapters is the means of achieving increased speed of research and development, and how the social and ethical dimension is affected through epistemological practices which partly have to enable the desired goals. So far we have seen the theoretical models and declaratory roles of translational

⁵⁷ I thank Torsten Wilholt for pointing out Crewe's review as a valuable source on AZT, which had changed my initial assessment of the case.

⁵⁸ See especially the work of Daniel Carpenter, in particular Carpenter, Zucker, and Avorn (2008) and recent publications as part of the FDA Project.

medicine, while we will now turn to how translations are actually done on the level of practice, as well as which limitations they face and which ethical considerations they give rise to. In order to explore that, it will be necessary to take a step back and look at two other approaches which comprise contemporary biomedical landscape and often intersect with translational medicine: personalized or precision medicine, and evidence-based medicine. A broader outlook is a prerequisite for understanding the epistemology of medicine and assessing the impact of translational medicine in particular. Personalized medicine is important because it provides the epistemic input for acceleration in both T1 and T2, while evidence-based medicine is important because it dominates the delivery of health care on the level of clinical assessment, in T2.

5. The landscape of biomedical research

5.1 Personalized or precision medicine

Personalized or precision medicine is an approach that tailors therapy to individual needs. It is often represented as “P4” medicine: predictive, preventive, personalized and participatory. The observations of highly variable drug responses have led to the development of a new scientific discipline from genetics, biochemistry, and pharmacology, namely pharmacogenetics, while advances in molecular medicine have led to pharmacogenomics which seeks to understand the molecular mechanisms of drug response (Vogenberg, Barash, and Pursel 2010, 560). In the new personalized approach patients’ gene variations guide the selection and dosage of drugs, and the aim is to minimize harmful side effects and achieve more successful outcomes (ibid.). Other expected benefits include indicating an individual’s susceptibility to certain diseases before their onset, allowing for monitoring and prevention. Unlike the blockbuster and generics one-size-fits-all model of prescribing, the idea is to make more effective clinical decisions for each patient (ibid.). Personalized medicine also takes into account individual lifestyle and environment, including diet, other inner and external exposures, as well as health history. It is closely related to systems biology, an approach based on modeling health and disease as emergent properties of a complex, dynamic, and multilevel biological system (Noell, Faner, and Agusti 2018, 1).

The integration of different omics research on the level of T1 enables more predictive outcomes of a therapy. It has already been said that omics technologies comprise research into human genome (genomics), mRNA (transcriptomics), proteins (proteomics) and metabolites

(metabolomics), but also research in emerging omics fields such as exposomics, which predicts individual disease risk related to the environment by focusing on the sum of all internal and external exposures (Canali 2016). The characteristic of omics research is that it approaches a specific biological sample in a non-targeted and non-reductionist manner, as a part of a unique and complex totality and a network of interactions. Investigations into such complex systems are facilitated by increasing amount of information and biological material stored in various databanks, such as population banks which store clinical data and lifestyle information. Biobanks, for example, are an especially important resource for personalized medicine. They are repositories of human tissue, blood, plasma, and other compounds that can be used as research material. For example, a rare type of tumorous tissue can be stored in a biobank in order to be used as research material for investigations into mutations that are specific to that type and its manifestation in a particular individual case. In personalized approaches the interaction between different factors is prioritized over isolated targets of research, such as individual genes, mutations or other variables of interest.

“The basic aspect of these approaches is that a complex system can be understood more thoroughly if considered as a whole. Systems biology and omics experiments differ from traditional studies, which are largely hypothesis driven or reductionist. By contrast, systems biology experiments are hypothesis-generating, using holistic approaches where no hypothesis is known or prescribed but all data are acquired and analysed to define a hypothesis that can be further tested.” (Horgan and Kenny 2011, 190)

For example, when breast cancer was first associated with specific genes, BRCA 1 and BRCA 2, this supported the view that single genes cause cancer. However, this hypothesis turned out to be wrong for most cancers, whose onset is associated with a number of factors. In most of the cases, single genetic factors only contribute to the likelihood of the occurrence of a disease, but the cause is usually associated with the interaction of many genes, as well as with the interaction of many biological levels (genes, molecules, tissues, organs, and organism) with inner and external exposures. Only by understanding a network of different interactions can a hypothesis about the causes of a specific disease in an individual patient be made, and the molecular pathways for intervention be adequately identified. The hypothesis-generating method is therefore data-driven, i.e. searches are made through a vast amount of information to find meaningful patterns, as opposed to accepting or rejecting previously generated hypotheses.

Besides the emphasis on prevention instead of intervention, optimizing therapy, and making drugs safer by avoiding adverse reactions, the added benefits of personalized medicine include increasing patients' compliance to treatments, reducing the time and cost of clinical trials, reviving drugs that failed early in trials or on the market (based on variations that make the drug effective and safe for a smaller group within the population), and reducing the overall cost of health care by decreasing the number of failed trials while increasing the number of successful outcomes (Vogenberg, Barash, and Pursel 2010, 561). By focusing on this approach producers can design more promising interventions and hope to recover the costs of even failed trials by targeting selected groups of patients that can still make use of their products. Regulators in this case have to focus on delineating the groups that benefit from a particular product from the ones that do not benefit from it. Most importantly, patients' health is likely to be improved by individualized therapies.

However, the cost of these treatments is usually very high and the long term effects are often not known. In 2017, the year of the two decade record approval rate, FDA approved more personalized medicines and tests compared to any year before (Bilkey et al. 2019, 2). Some of these treatments were biomarker specific therapies reliant on genetic testing. Three approvals were the first gene therapies ever approved by the FDA, including voretigene neparvovec (Luxturna) for retinal dystrophy, the first to treat an inherited disease. With a price of 425.000 dollars per eye, it is the most expensive medicine in the USA per dose (ibid.).

In December 2016 nusinersen (Spinraza) was approved for early spinal muscular atrophy, a condition that affects approximately one out of 10.000 people at birth (ibid.). The case of Spinraza has gained a lot of publicity as it got accelerated approval through managed access agreements⁵⁹ and as countries have gradually started making it available with different restrictions concerning age and clinical picture. For example, in Australia Spinraza was listed on the Pharmaceuticals Benefits Scheme from June 2018 for patients younger than 18, while in Britain it was considered too expensive and its long-term effectiveness too uncertain for inclusion in the National Health Service, until an agreement with the producer was finally reached in July 2019 (Bilkey et al. 2019, 2; Muscular Dystrophy UK: "Spinraza", n.d.). In early 2018 Croatia has been one of the first countries in Europe to make Spinraza available with

⁵⁹ Managed access agreement allows patients to access a treatment prior to a full regulatory approval, so that data about effectiveness can be gathered while ensuring access to selected groups of patients who are thought to benefit the most from it (Muscular Dystrophy UK: "Spinraza", n.d.).

reimbursement, but only for patients younger than 18 and who are not on a respirator. In the course of the last two years patient groups have regularly been protesting and pressuring government to put the medication on the list of the Croatian Health Insurance Fund with no restrictions when it comes to age and patient's condition (Jureško 2019a), which was finally achieved in July 2019, but only for treatment in the first six months (Jureško 2019b). Spinraza, however, has to be delivered for the remainder of patient's life, and it costs 125.000 dollars per injection, meaning 750.000 dollars in the first year of treatment per patient, and 375.000 dollars for each subsequent year (Bilkey et al. 2019, 3).

Due to the increased cost which decreases accessibility, coupled with unknown long term effects which give rise to regulatory issues, personalized medicine faces substantive challenges. Other ethical, social, and legal issues arise from concerns about patients' privacy and the confidentiality of health information, since databases have to store large amounts of patient data, as well as from concerns related to overdiagnosing for issues that do not need to be treated or that will not develop into a disease.

5.1.1 The relationship between personalized medicine and translational medicine

The relationship between translational and personalized medicine is very close. Miriam Solomon discusses a case of translational success, a novel therapy for lymphocytic leukemia based on modified T cells, which can also be described as a success of personalized medicine (Solomon 2015, 159-160). In fact, the treatment was advertised as an example of personalized and not translational medicine, which she explains by the fact that patients and their families are more likely to be attracted to the idea of personalized medicine than to the idea of translational medicine (*ibid.*, p. 160, footnote 13). I contend that a reason for this is that "personalization" signals care for the patient, while "translation" signals a path to a product, as suggested in "bench to market" depiction. Another possibility is that "translation" is simply a vague term, not distinctive and understandable enough. The most likely, closely related reason is that translational medicine boils down to the aim of translating, while personalized medicine is more informative about its method. For Solomon, the case of therapy based on modified T cells is a clear case of translational success because it fulfils the T1 criterion. According to the T1 criterion, any product that reaches in-human trials from animal studies and shows efficacy has been successfully T1 translated. It is thus not surprising that all personalized medicine successes are also translational successes. Let me explain this in more detail.

Successfully bridging the gap in T1 means being T1 translated – a therapy has reached in-human studies and showed efficacy against a disease after assessment in preclinical studies. Bridging every further translational gap will be a new translational success, so we will hopefully have T2, T3, T4, and T5 translations. This means that we can have successful translations even though a therapy may fail in phase III trial, which is what happens to an estimated 70-75% drugs in the USA (FDA 2018c). Since each and every translation is a particular goal, translational results are made more obvious when particularized to sub-steps.

In the case of personalized medicine, overall translation will occur faster because certain translations will often not be necessary, namely phase III trials. They are usually conducted to assess efficacy of a treatment in a large number of patients, hence indicating effectiveness in non-experimental settings and in broader population of diseased patients. Failures in phase III can arise from a lack of efficacy, issues with safety, lack of funding to complete a trial, failing to maintain good manufacturing protocols or to follow regulatory guidance, or due to problems with patient recruitment, enrollment, and retention (Fogel 2018, 156). However, when therapies are individualized, hence designed for individuals and groups with a particular genetic setup or suffering from a particular genetic disease which can be more or less rare, the usual phase III trials will often be either impossible (in the case of very rare diseases), either superfluous (since it is expected that they work for some and not for most), or even unethical, if they unjustifiably extend the time for a life-saving or otherwise high-stake therapy to become available.

As a consequence of that, personalized medicine has led to a development in the design of early-stage clinical trials, precisely in order to improve access and evaluate efficacy earlier and more efficiently. Some of the novelties include the emergence of adaptive design trials (Garralda et al. 2019, 549). The Ebola ring trial design, for example, is a case of adaptive trial because it uses ring clusters instead of individualized randomization, though it is not a case of a personalized therapy. Some examples of adaptive measures in clinical trials include early stopping rules in the case of a lack of efficacy or in the case of unacceptable toxicity, as well as changing doses or drug schedules. An example of a novel adaptation strategy is the use of accrual design: after the initial ‘learning phase’, in the ‘adaptive phase’ the ratio of patients randomly assigned to the experimental arm as opposed to the control arm changes from the standard 1:1 to increase the proportion of patients in the arm that is performing better, which increases the statistical power to detect clinical benefit (ibid., p. 551). Randomization to different drugs or combinations of drugs can also be changed based on results generated in real time. The goal of adaptive designs is to learn from the data in the early trials and apply the

knowledge as soon as possible. The modifications to the study design can include eligibility criteria, proportion of randomization, and seamless design (ibid.)

Adaptive enrichment is a term that refers to the modification of the patient eligibility criteria. If analysis shows that one subgroup has a more favorable response, the trial can be “enriched” by modifying it to either exclusively or predominantly enroll patients from this subgroup (Thorlund et al. 2018, 2). Response adaptive randomization allows for changes in the randomization ratio during the trial, so that the newly enrolled patients can be assigned to the treatment arm, as it is the case in the accrual design (ibid.). Furthermore, seamless adaptive trial design allows for proceeding from phase II to phase III trial in a non-standard way. The results from the phase II trial are used to determine the initial patient allocation ratio, the planned total sample size (which can be rather smaller than the usual phase III samples that normally include from 300 to several thousand patients), and a potentially enriched set of patients, i.e. patients that are thought to benefit the most from the intervention (ibid.).

Furthermore, machine learning methods have been developed to help guide the best matched targeted therapy. In some adaptive design studies predictive algorithms incorporate prior knowledge based on computer models of drug sensitivity, biotechnological experiments, preclinical or early clinical data, and search for the best match (Garraida et al. 2019, 552).

“This may be useful when multiple druggable alterations are identified in a patient’s tumour sample and more than one agent is available for testing; and when one driver genomic event is identified, and the investigator has to select among various drugs with overlapping mechanisms of action (targeting the same driver event) but with different potency/activity according to coexisting genomic alterations. These ‘machine-learning predictive models’ can complement molecular tumour boards efforts to identify the ‘best guess’.” (ibid.)

The problem with adaptive trials is that results can be difficult to interpret due to biases in research design, especially towards favoring false positives. Adaptive trials research can be blinded or non-blinded, but it usually includes smaller sample sizes, so there is an increased risks for misleading statistical significance results. Furthermore, each trial is adapted in a particular way, so informed consent and efficient communicating of risks and benefits can also pose a problem (ibid.). The trials vary over time, include different adaptive measures, so the claims of efficacy can be very uncertain, since confounding factors and biases cannot be adequately canceled out.

Contemporary translations are very likely to occur on the terrain of personalized therapies and they occur there faster due to changes in the drug discovery methods and clinical assessment route. The relation is this: all personalized medicine successes are translational successes, but not all translational successes are personalized medicine successes. This is so because translation is the aim, and precision research is one of the means to achieve it. It delivers on the requirement of speed most successfully, and it is to a large degree enabled by biobanks and omics advances developed through private-public hubs and partnerships which have been created or incentivized as a part of translational initiatives. In fact, personalized medicine is an approach highly represented in the Roadmap programs through an emphasis on omics research, information technologies, and biobanks, and it aligns especially well with the translational aim of accelerating health care delivery.

What also connects the two approaches is the fact that both translational and personalized medicine have gained widespread recognition and incentive through ambitious government initiatives. Translational medicine gained attention with the NIH Roadmap in 2004, while Precision Medicine Initiative originated with Barack Obama's 2015 State of the Union Address. The initiative was launched with 215 million dollar investment in the US 2016 budget "to pioneer a new model of patient-powered research that promises to accelerate biomedical discoveries and provide clinicians with new tools, knowledge, and therapies to select which treatments will work best for which patients" (The White House 2015). Justification in terms of patient empowerment and the acceleration of discovery and research is shared in both initiatives, only now the focus is on individualized therapies. There are good reasons to think that personalized medicine is the most prominent and tangible success of translational initiatives, especially if the criteria are: demonstrated efficacy of a therapy, efficiency in bridging translational gaps, acceleration towards implementation, and finally, actual implementation in practice.

But implementation does not grant access, as we have seen from the case of Spinraza. Implementation in the form of regulatory approval and established clinical guidelines is only a necessary, but not a sufficient condition for access. It may or may not correlate with people actually being treated. Furthermore, even if they are being treated, the effectiveness of many treatments has yet to be comprehensively assessed. The era of individualized therapies has only recently began and a number of uncertainties related to effectiveness, safety and privacy pose a challenge from an ethical and a regulatory perspective.

Translational medicine is a cluster of stage transitions in the development of a medical product at the intersection of basic and clinical research, and more broadly – prevention, guidelines and health policy. Personalized medicine is a way of doing medical research by using big data, advances in omics research, information technology, high throughput screening technologies, and machine learning to bring about medical innovation tailored to individual needs of the patient. Since both translational and personalized medicine are dedicated to making health care delivery faster and more efficient, their successful cooperation is not a surprise. Quite the opposite – it was the biobanks collected as a part of translational initiatives in the early 2000s that have made it possible to personalize medicine in the 2010s. Notably, both initiatives highly value speed in discovery, research, and development, which is not only a success of science, but of a larger cooperative work and exchange of many stakeholders, institutions and disciplinary cultures. In this case, insistence on speed is motivated by high social stakes, primarily in improving the health of the population, but also the cost-effectiveness of medical therapies.

5.2 Evidence-based medicine

Evidence-based medicine (EBM) is an epistemological approach and practice in clinical medicine that has dominated medical decision making since the early 1990s. The epistemological claim of EBM is that medical knowledge should be based on the best possible scientific evidence and that best evidence relies on epidemiological and biostatistical methods (Solomon 2011, 2015). David Sackett, one of the pioneers of evidence-based medicine, defines it as:

“(…) the conscientious, explicit, and judicious use of current best evidence in making decisions about the care of individual patients. The practice of evidence-based medicine means integrating individual clinical expertise with the best available evidence from systematic research.” (Sackett et al. 1996, 71)

The availability of large online depositories of clinical trials contributed to the spreading of the emerging medical movement, since the Internet has made such evidence broadly accessible and continually updated. Up to then, medical practitioners were not used to engage in thorough searches of this material, but were rather making clinical judgments on the basis of established protocols and expert knowledge. EBM is characterized by the hierarchy of evidence on whose

top are meta-analyses, followed by systemic reviews and randomized controlled trials (RCTs) as the “gold standard” evidence, especially double or triple blinded trials. Meta-analysis integrates data from different trials on the same issue to give an overall single statistical result. A systematic review is a search of the literature which compares similar clinical trials and evaluates them. The evidence from these sources, especially when they are well-designed and well-conducted, is considered to be most reliable in avoiding biases such as selection bias and confirmation bias, as well as in cancelling out confounding factors that might distort judgments about safety and efficacy. On the lowest levels of evidence hierarchy are case reports, observational studies, expert judgments, and evidence of mechanisms which underlie interventions. EBM originated with clinical epidemiologists at McMaster University and Oxford University in the 1970s and 1980s, and developed into a new paradigm in the early 1990s through the Evidence-Based Medicine Working Group (Solomon 2011, 452; 2015, 105-111). It has been embraced by medical training programs and leading medical journals in the world as an established theory and practice of clinical medicine and clinical decision making. Nowadays, evidence-based practice is an approach used not only to inform decision making in medical contexts, but also in public health, nursing, management, social policy, and library science (Jukola 2019, 1).

Solomon (2011, 2015) emphasizes that EBM nonetheless has problems with biases, especially with publication bias and pharmaceutical funding bias. Publication bias occurs when studies with null, i.e. negative results are not accepted for publication because they are thought to be less significant than positive results. This can keep valuable research out of sight and distort the balance of findings. Publication bias can also motivate unacceptable research practices, like adjusting research design in order to bring about positive results. Pharmaceutical funding bias comes from the fact that trials are mostly funded and executed by pharmaceutical companies which are profit driven and therefore bias prone, i.e. they can subtly influence the design and evaluation of therapies. Such intrusion of biases questions the internal and external validity of EBM’s claims.

The controversies around EBM have also arisen out of the “hegemonic” nature of the movement.⁶⁰ Concerns have been raised as to whether EBM is really integrating different kinds of evidence, rather than downplaying valuable evidence on the lower levels of evidence hierarchy, especially clinical expert judgment (Tonelli 1998) and pathophysiological reasoning

⁶⁰ “Intellectual hegemony” (Berwick 2005, 315)

about the causes of diseases and the underlying mechanisms of intervention (Russo and Williamson 2007). It has also been asserted that EBM methods favor populations rather than the individual as the primary unit of care. Since evidence from trials gives guidance based on average and not individual responses, effectiveness in individual cases cannot be predicted in a warranted way (Cartwright 2011, 2012). This brings in question the external validity of EBM's best evidence.

Scholars and medical practitioners have also emphasized that each clinical problem should be assessed based on evidence which provides best guidance for a particular case, and not based on general "cookbook medicine" guidance (Knaapen 2014). Especially, that clinical judgment is indispensable for bringing about sound medical decisions (Tonelli 1998). A multitude of particularities and variabilities has to be taken into account, which is not possible to incorporate when relying on statistics only. In a word, the most prominent criticism of EBM centers around the problems of internal and external validity of the knowledge produced by EBM's best standards, as well as around the overlooked role of clinical experience, expert judgment, patient variability, patients' goals and values, local health care constraints in conducting large scale RCTs, and basic medical sciences (Solomon 2015, 116, 133-154).

5.2.1 The relationship between evidence-based medicine and translational medicine

I will focus exclusively on Miriam Solomon's treatment of the last critical point, namely the absence of basic science and causal reasoning in EBM, since she is the only philosopher of science to bring this in connection with the emergence of translational medicine. She is in fact the only philosopher (as far as I know) who has included translational medicine in any systematic overview of contemporary biomedical practices, and she has done this in her 2015 book *Making Medical Knowledge*. I will follow her account in order to keep the focus on EBM insofar as it gives a perspective for understanding and discussing translational medicine, and more broadly, trade-offs connected to acceleration of research on the clinical level. Moreover, the only subsequent philosophical work that directly addresses translational medicine does so in response to Solomon, and it is Mark Robinson's 2019 article "Financializing epistemic norms in contemporary biomedical innovation". I will say more about Robinson's contribution in the conclusion of part II because he makes explicit the relation between epistemology and the

broader socio-economic context in which translational medicine is done, which will serve as a transition to the values discussion in part III.

Miriam Solomon is explicit in linking the emergence of translational medicine with three distinct factors:

“First, there were wildly optimistic expectations of the basic science work of the Human Genome Project; second, there was lack of recognition of the incompleteness of evidence-based medicine methodology; and third, there was ignorance about what it takes to disseminate new practices.” (2015, 177)

I mostly agree with Solomon on this, especially when it comes to the first and the third point. Clinical applications of the Human Genome Project were enthusiastically expected, and translational initiatives have set up a framework to, indeed, translate them to what we know today as personalized medicine. This translational effort has been centered on bridging the T1 gap, from laboratory to the clinic. Implementation and dissemination, on the other hand, are at the heart of the T2’ gap broadly construed. This gap has also been acknowledged in translational medicine, and the collaboration of multiple stakeholders has been encouraged and incentivized in order to overcome it. Physical proximity of laboratories and clinical centers; better flow of information about new therapies; direct engagement with the patients inside and outside of clinics; increased awareness of the importance of guidelines, recommendations, and preventive practices; outreach to the public through various initiatives; continuous education and involvement of general practitioners – these are measures to address T2’, i.e. to address the perceived ignorance about dissemination (and implementation) put forward in Solomon’s third reason for the emergence of translational medicine.

In this section I will focus on Solomon’s second reason for the emergence of translational medicine, namely the recognition of the incompleteness of evidence-based medicine methodology. Solomon’s approach suggests that translational medicine is an improvement on EBM or some kind of an alternative to EBM, which is not the case. They co-exist complementarily and cover different areas of medical research, as well as different goals and methods to achieve them. In her 2011 article Solomon does not clearly acknowledge this, though she acknowledges it in her 2015 book. I will now explain this in more detail.

In her book, Solomon focuses on a particular line of EBM criticism, namely the apparent lack of basic science in EBM methods. Evidence from clinical trials is supposed to show whether

interventions are effective⁶¹ and not why or how they are effective. The judgment about effectiveness is made when there is a correlation between the intervention and a positive outcome, so the resulting knowledge is about effects and not about mechanisms. Mechanistic causal evidence is included in the evidence hierarchy, but it is usually on the lower levels of the pyramid, alongside anecdotal evidence and expert consensus. The role of mechanistic evidence in EBM has been widely discussed in philosophy of science, and Solomon argues against a shared assumption in the work of Howick (2011), Andersen (2012), and Clarke et al. (2013).

The assumption is that any instance of pathophysiological or mechanistic reasoning should be understood as mechanistic evidence that finds its place in the evidence hierarchy and needs to be integrated with other kinds of evidence. For Solomon, mechanistic reasoning is reasoning about the causes which provides understanding of the pathophysiological processes, while mechanistic evidence is evidence that a certain intervention will work based on underlying mechanisms. Solomon challenges the assumption that mechanistic reasoning is the same as mechanistic evidence, and that its role thus exhausts somewhere in the evidence hierarchy. She argues that “mechanistic reasoning provides weak evidence at best, but it has important non-evidential roles” (Solomon 2015, 123-124), and that using mechanistic evidence and mechanistic reasoning interchangeably leads to confusion.

“To be sure, we can have evidence for mechanisms, but that is evidence that the mechanisms operate, not evidence that a particular proposed intervention (which depends on more than the hypothesized mechanisms, even if those mechanisms exist) will work. We could have strong evidence that the mechanisms operate, yet no evidence (or the weakest of evidence) that a particular proposed therapy will have the desired effect.” (Solomon 2015, 123)

To illustrate her claim, she uses Howick’s (2011) example about the false hypothesis that hormone replacement therapy reduces cardiac mortality. Although there was strong evidence of hormonal effects on blood lipids, there was no evidence that hormonal replacement therapy would lead to reduced cardiac mortality, since the knowledge of relevant mechanisms was lacking. The evidence was exaggerated in favor of the initial hypothesis, which turned out to be wrong. According to Solomon, this was a case of the evidence of mechanisms, but not

⁶¹ She uses “effectiveness” but I prefer “efficacy”, as long as we are talking about clinical trials and not non-experimental settings. But literature is also divided on this point and we can find both usages when talking about stage III clinical trials.

mechanistic evidence that the intervention will work. Mechanistic evidence about the effectiveness of the intervention was very weak and it should not have been considered relevant evidence. But, and this is now Solomon's point of departure from Howick and others – mechanistic reasoning plays other roles than the evidential role. It is a tool for discovery. Pathophysiological or mechanistic reasoning is responsible for bringing a therapy to the clinical trials in the first place.

“So evidence-based medicine should not discount mechanistic reasoning unless it want to bite the hand that feeds it!” (Solomon 2015, 125)

Furthermore, without incorporating a role for mechanistic reasoning apart from its evidential role, EBM is “not a complete epistemology of medicine” (ibid., p. 121).

“What is at stake with whether or not we value mechanistic reasoning for its evidential role or for some other role? I think what is at stake is whether or not evidence-based medicine is a complete epistemology of medicine.” (ibid., p. 124)

This only makes sense if EBM was indeed conceived and presented as a complete epistemology of medicine, which is not the case. EBM is about evidence integration for clinical decision making and about favoring randomized controlled trials and meta-studies as a source of high quality evidence, and not about basic research and technological innovation development. Solomon, however, emphasizes the importance of translational medicine in incorporating mechanistic causal reasoning which is not found in EBM. The “tool for discovery” role for mechanistic reasoning as opposed to the role of trials in effectiveness assessment during clinical research, she argues, corresponds to the “romantic view of science” (ibid., p. 125) divided into the “context of discovery” and the “context of justification” as put forward by Hans Reichenbach in 1938. In the process of making medical knowledge, mechanistic reasoning is a part of the context of discovery, while clinical trials are a part of the context of justification.

“The ‘context of discovery’ allows any creative methods but the ‘context of justification’ is the place for rigor in evaluating the creative ideas developed in the context of justification.” (ibid., p. 125)

Interestingly, Solomon identifies translational medicine with the basic end of its bridging role (understanding mechanisms), while we could see that the primary drive for its formalization was a call for more applied research, and that it has been criticized for this imperative and the appropriate means to address it (Maienschein et al. 2008; Fang and Casadevall 2010).

Discovery is what supposedly happens at the bench side, and translation is the process of bringing it to the bedside. In Solomon's account translation gets its primary role at the bench side again.

In order to better understand Solomon's view, some context for her argument is needed. It is first important to emphasize that EBM has been extremely influential, revolutionary, and fairly successful, and that it still is the established way of doing clinical medicine. As such, it has gained a lot of philosophical attention, and the emergence of translational medicine is almost necessarily evaluated in the light of EBM, since it can be expected that the strengths of the movement on the rise would be at least partly addressing the weaknesses of the movement that is currently prominent.⁶² In her book Solomon sets out to compare and evaluate several ways of doing medicine that are on the table: EBM, translational medicine, and narrative medicine, with a special interest in consensus conferences which are in the realm of dissemination. Personalized medicine is not included in her discussion though she recognizes its prominence and soon rise.⁶³ Solomon's conclusion is that there is an "untidy methodological pluralism" (2015, 224) in doing medical research. Untidy means that there is no hierarchy or linearity in the medical approaches she discusses. They exist in parallel and each has its merits and downsides. Still, there is an important quality of this pluralism:

"Much of the time, the methods do not compete with each other; they each have roles to play, often at different stages of research." (ibid., p. 228)

⁶² At least in the vein of Thomas Kuhn's (1962) scientific change account, which Solomon takes up when discussing EBM as a Kuhnian paradigm (2011).

⁶³ "Is 'personalized medicine' the next trend in epistemology of medicine? It has already achieved some prominence. I am not yet convinced that there is enough in the way of new methods, or even reclaimed old methods, to consider personalized medicine a new way of doing medical research and/or practice. Perhaps such methods will come in due course." (Solomon 2015, 228)

Contrary to this view, I find much more scientific change in personalized medicine as compared to EBM, than in translational medicine. But the tides can nowadays change in the course of five years when it comes to state of the art medicine and prevalent medical movements. It should not be forgotten that the record FDA approval rates with significant increase of personalized medicine treatments occurred only very recently, in 2017 and 2018.

I would like to underline this point and put things into perspective based on what we know about translational medicine, personalized medicine, and the merits of evidence-based medicine.

The development of adaptive trials in clinical research can be seen as a response to the limitations of RCTs as the gold standard in EBM. Though RCTs are assumed to provide best quality evidence, they are also considered inefficient in terms of cost and length, lacking the essential translational quality of actually pulling the innovation through towards approval and implementation. But the actual change in practice towards a broader use of adaptive trial design occurred later, with personalized medicine, and is fraught with risks and uncertainties. Fairly enough, translational efforts have contributed.

Solomon, however, puts forward the argument that what has been missing from EBM is mechanistic reasoning, the basic science quality, only to be found in the emerging translational approach. But basic science did not disappear in the EBM era, it was indeed providing therapies for clinical research, so if translational initiatives emphasize it, they do not emphasize anything new. And if basic science is at the heart of translational medicine, it cannot be praised as the successor of EBM, since they cover fundamentally different areas. Even though Solomon recognizes distinct areas of EBM and translational medicine, she sometimes treats translational medicine as a moderate improvement on evidence based medicine, grounded on the finally recognized non-evidential role for mechanisms that EBM has neglected. This was especially prominent in her earlier work from 2011:

“There is some indication that EBM is now past its peak, and being overshadowed in part by a new approach, that of ‘translational medicine’.” (2011, 453)

“With the recent emphasis on translational medicine, we are seeing a restoration of the recognition that clinical research requires an engagement with basic theory (e.g. physiological, genetic, biochemical) and a range of empirical techniques such as bedside observation, laboratory and animal studies.” (ibid., p. 464)

Translational medicine cannot be an improvement on EBM since they are about different things: translational medicine is a theory and practice of technology innovation, while EBM is a theory and practice of best clinical evidence and how to use it. What is especially misleading is that Solomon settles for an untidy pluralism of medical methodologies in her 2015 book, when each of the medicine movements she discusses is about a different stage of medical research and practice. Their difference is on a more fundamental level than “often they don’t

compete” or even “much of the time”. There is a clearer difference between the area that each covers and the goals that they pertain to. Jonathan Fuller (2016) shares this view:

“In other words, translational medicine applies to medical research, consensus conferences apply to knowledge dissemination, and EBM and narrative medicine apply to clinical practice. The main purpose towards which each method is put is unique: translational medicine develops new medical technologies, consensus conferences develop consensus statements or clinical guidelines (often pertaining to those technologies), EBM appraises evidence and applies it in clinical practice, and narrative medicine uses narrative techniques at the bedside. Thus, their domains are less overlapping and arranged more linearly than Solomon’s untidy pluralism might suggest.”

However, I do not think that pluralism is a completely wrong way of conceptualizing contemporary biomedical research. What is misleading is the egalitarian view of medical movements that Solomon suggests, when in fact they are targeted to different stages of medical research and practice, i.e. they are more specialized. But inside each of them, there is a pluralism of methods to reach the envisaged goals and find room for improvement. For example, a multidisciplinary group around EBM+ project is seeking to improve the ways in which evidence-based medicine handles evidence of mechanisms (see “EBM+”, n.d.). This initiative is trying to improve or provide alternative to the ‘orthodox’ EBM inside of EBM itself.⁶⁴ Another project supporting EBM+ is CauseHealth, or Causation, Complexity, and Evidence in Health Sciences, initiated by philosophers Rani Lill Anjum and Stephen Mumford, which explores multifactorial causation in medicine from a dispositionalist point of view. Dispositionalist account explains how the dispositions of the recipient of an intervention contribute to the produced effect. To be fair, these are primarily philosophical projects, but they collaborate with medical practitioners and share an assumption that EBM should be improved, broadened and changed, but rather from within than from the outside.

⁶⁴ Solomon does not mention at all the influential paper by Russo and Williamson (2007) about the necessity of both probabilistic and mechanistic evidence to establish a causal claim in medicine. Fairly enough, she is not convinced by the evidential role of mechanisms. Still, there seems to be more to EBM than the simplistic statistical analysis, which is recognized inside the movement itself. Williamson is one of the founders of EBM+.

Solomon finds support for her account of the emergence of translational medicine in Berwick (2005) who also points out some limitations of EBM. His paper invites for more recognition of the excluded methods and suggests to broaden EBM in the direction of more “pragmatic science” by:

“- tracking effects over time, especially with graphs (rather than summarizing with statistics that do not retain the information involved in sequences);

- using local knowledge – the knowledge of local workers – in measurement (rather than relegating measurement to people least familiar with the subject matter and work);

- integrating detailed process knowledge into the work of interpretation (inviting observers to comment on what they notice rather than “blinding” them to protect them against what they know);

- using small samples and short experimental cycles to learn quickly (rather than overpowering studies and delaying new theories with samples larger than needed at the time); and

- employing powerful multifactorial designs (rather than univariate ones when the better questions for the time are formative, not summative).” (Berwick 2005, 316)

Berwick goes on to conclude that “pragmatic science of this type is alive and well” (ibid.) and Solomon is ready to name it – translational medicine:

“Berwick did not use the term ‘translational medicine’ – perhaps the term was not yet in wide enough usage – but it is clear from the context that ‘pragmatic science’ involves the same kind of observations and trial and error experimentation as does translational science.” (Solomon 2015, 170)

I agree that what Berwick is referring to is best captured by translational efforts, but clearly on the level of clinical trial design and not on the level of basic research that Solomon puts forward as the most significant asset of translational medicine. It is also worth reminding that translational medicine is included in the NIH Roadmap as a part of “Reengineering the clinical research enterprise” cluster of initiatives, though it has always been represented as a cluster of multiple translations in development of a product. In any case, Berwick is referring to clinical trials: smaller samples (of patients), learning quickly (in phases I and II), multifactorial design (of trials), integrating process knowledge (and adapting the design) – all of this is supposed to

accelerate research and make it more “pragmatic”, as opposed to EBM standard which supposedly provides high quality evidence but is long and inefficient, since most of the drugs fail in phase III trials, and often the effectivity of even the ones that do not fail is questioned. Berwick’s pragmatic approach maps well onto the adaptive trial design, which only came in focus later, with personalized medicine.

The relation between RCTs and adaptive trials involves a trade-off: a more reliable and unbiased method is long, especially for cases where urgent action is needed, while the adaptive method is fast, but not as reliable. In personalized medicine better outcomes are expected from interventions based on data-driven hypotheses that take into account patients’ individual molecular and other data, and are thus more likely to result in positive outcomes. But in the cases of traditional drug design, such as in Ebola ca suffit! trial and AZT trials where adaptive measures were also applied, we can recognize a tension between speed and reliability.

5.3 Conclusion

I have argued that personalized medicine is a success story of translational efforts achieved through means that are both internal and external to scientific research, although there are doubts as to whether the distinction can be maintained. Why is this so? Personalized medicine is a science-society based way of achieving translational goals in a systematic and predictable fashion that traditional approaches have been short of. While science provides methods, the expanded material base of biological, molecular, and life-style data has to be collectively gathered and continuously updated in order to be useful. The integrative achievements that lie in the background of personalized approaches have to be institutionally and interactively supported by various scientific and non-scientific parties. The external measures of translational efforts aiming at better connection of academy and industry have contributed significantly to the uptake of practices and development of advances related to personalized medicine, and to a better recognition of systems biology as its theoretical background. Systems biology is the theory of a holistic approach to organism, while personalized medicine is systems biology in practice, backed up by technological advances and a broadened material basis of science known as big data.

The success of personalized medicine is especially prominent if we take that its primary criterion is accelerating discovery and research towards implementable products and practices, especially on the level of T1. Material prerequisites and integrative practices that have led to its

emergence have been highly financially and politically incentivized in order to bring about scientific and social change. The renewed academy–industry–society relation that has contributed to recent scientific advances makes them not only scientifically successful, but also economically, socially, and politically successful. For example, databanks are collections of diverse information which can, but does not have to become evidence. In order to become evidence, various data sets first have to be collected, then treated, stored, distributed, and used in particular ways (see Leonelli 2016). Since integration is the key to understanding in holistic approaches, it has to be facilitated. The integration of data is facilitated by a complex interaction of scientific and non-scientific factors, i.e. internal and external measures.

EBM is not an alternative to personalized medicine in the T1 stage, since the focus of EBM is on clinical practice and favoring particular ways of doing clinical research, namely RCTs. However, personalized medicine has resources to challenge the dominance of EBM's standards on the level of clinical research by the practice of adaptive trials. They were not as developed when translational research initiatives were emerging, but were nonetheless encouraged by them. EBM's high standards of evidence do not align well with the goal of accelerating research that translational medicine emphasizes. RCTs are more reliable the longer they take and the larger the samples are, granted that randomization, blinding, and assessing the outcomes is performed well. However, according to Stegenga (2018) even the best RCT standards do not make it likely that the interventions will be effective for a larger population of patients. Nonetheless, RCTs remain the most reliable standard for determining the effectiveness of medical interventions because they aim at reducing the influence of confounders and biases.

However, there are arguments in favor of a faster uptake of new medications in clinical trials. First, because of social, non-epistemic reasons, such as those presented in the case of Ebola and AIDS epidemics. Second, because of epistemic reasons, such as those coming from better understanding of interactions on the molecular level coupled with favorable systems effects recognized early on, which is an asset of adaptive trials in personalized medicine. Better understanding of mechanisms can make RCTs redundant, especially if therapies are designed for a restricted number of users.

We can read Solomon's criticism of EBM along these lines (understanding mechanisms as opposed to statistical methods of high quality evidence in EBM), but she is not making explicit the transition towards personalized therapies which are aimed at a narrower patient group, as opposed to ideally widely generalizable results of RCTs. Since this important dimension is

missing from her account, and the relation between translational and personalized medicine is thus underdeveloped, it is somewhat hard to get a comprehensive picture of the pluralism of medical movements that she suggests. On the other hand, by uttering the criteria of acceleration that I advance, these relations become much more transparent. I ultimately agree with Solomon that some aspects of translational medicine are a reaction to EBM, but these aspects are almost exclusively the ones pertaining to adaptive trials as used in personalized medicine. In this case, translational label is superfluous and only causes confusion, which is indirectly confirmed by its notable absence from philosophical discussions.

Adaptive trials are thus said to be prioritized over RCTs under one or both of the following conditions: that an action needs to be urgently taken, or that the molecular base of drug action is well understood, as well as systems effects carefully assessed, which would make the overall outcome more predictive for a certain number of cases. But even when both the criteria of urgency and of better predictive capacity are satisfied, the measures towards acceleration are fraught with uncertainties, since long-term effects cannot be adequately assessed regardless of how urgent the action is and how well established the mechanism is. This makes clear that the trade-offs between different goals and values are inevitable, of which more will be said in Chapter 7 and in Part III. Furthermore, it is important to identify some inherent limitations to the acceleration of biomedical research, to which we turn next.

6. Limitations to the acceleration of biomedical research

I have argued so far that the accelerated contemporary translations are mostly happening in the personalized medicine context. This, however, does not mean that successful translations cannot be happening outside of the personalized context, it only means that they cannot be systematically accelerated by any scientific measure. To a lesser degree this is also true of personalized medicine, especially since the internal/external distinction collapses in the contemporary biomedical settings. In the next section I will focus on two epistemological problems in preclinical and clinical stages of translational process which present limitations for accelerating research. The first has to do with strategies in early drug discovery and the second with measuring the effectiveness of medical interventions. I will also present measures that address the former of the two problems.

6.1 The case of early drug discovery

Methods in early drug discovery are broadly based either on the complexity of therapeutic interventions to target specific biological targets or on the demonstrated system effects. Both approaches have been considerably successful, but their successful integration seems to be a matter of coincidence or a lucky guess, as argued by Mathias Adam (2011). This kind of shortcoming cannot provide for a reasonable accelerating of the research process, unless in terms of optimizing the chances for coincidences, lucky on both effectiveness and safety. The rate of serendipitous coincidences, however, can be significantly increased by setting an institutional framework that fosters their occurrence, which corresponds to what we find in translational, and especially personalized medicine. Translational medicine, most broadly construed to account for both personalized and non-personalized approaches, is a systemized attempt to foster serendipitous discoveries.

In order to argue for this, I will start again from Solomon's characterization of translational medicine:

“Translational medicine (T1, the main part of translational medicine) involves causal reasoning and informal experimentation, using hypothesis about the mechanisms involved and trial and error in interventions.” (Solomon 2015, 226)

“Trial and error in interventions” can have four meanings here. First, Solomon might be using trial and error experimentation to refer to adaptive design in early stage trials which allows for changes in doses, variants of drugs, and randomization based on the learning phase. The experimenters can be thus said to try, to err, and then to change something and try again. This happens on the level of early clinical research after the molecular base of the intervention has been hypothesized. The opposition to this are strict evidence-based guidelines which ensure that the research is randomized and blinded in order to produce reliable evidence.

Usually however, trial and error experimentation is the traditional empirical method of early drug discovery. It is older than rational drug design (the discovery method of blockbuster drugs in the 1970s and 1980s) and high-throughput screening (automatized random search), and largely reliant on serendipitous findings. Traditional trial and error experimentation focuses on system effects, i.e. the response of the whole organism, and not the response of specific targets treated by rationally designed drugs that intervene on molecules in isolation, or genetic interventions that intervene on individual genes. Traditional, historical empirical search

strategies involve randomly chosen substances or chemical modifications of existing drugs that are then tested in animal models. They are identified empirically on the basis of observable systems effects, independent from a scientific understanding of the underlying molecular interventions (Adam 2011, 68).

Third, Solomon might be interpreting high-throughput screening technologies that aid personalized medicine on the level of drug discovery as a form of trial and error experimentation. High-throughput screening is a process by which large numbers of compounds can be tested automatically for activity as inhibitors (antagonists) or activators (agonists) of a particular biological target, with the goal to identify high-quality 'hits' or 'leads' (compounds that affect the target in desired manner) that are active at a fairly low concentration and have a new structure (Broach and Thorner 1996, 14). Since the early 1990s, high-throughput screening was developed as an efficient method to test a large number of substances empirically (Adam 2011, 68). It involves a robotized, highly efficient hit and miss strategy, which can be described as trial and error since it optimizes a chance for a hit, but what is characteristic of the high-throughput method is that it is concerned with local, reductive problems, i.e. with the activity of drug candidates on their targets. This type of reductionism is not traditional at all, since the traditional trial and error experimentation discovers efficacious compounds based on their holistic effects, i.e. effects on the level of the organism. However, when combined with the methods of systems biology, high-throughput screening facilitates the search for molecular pathways of intervention on the local level, since understanding of the whole network of interactions still has to be combined with some local, narrowly targeted action in order to efficiently intervene. High-throughput facilitates the search for the site of local action, while systems approach accounts for the system response, so the overall approach is not reductivist, but combined and integrated.

The practice of combining these two approaches, high-throughput screening and systems biology, is the fourth possible meaning of “trial and error in interventions” as used by Solomon. She might be referring to a “rejuvenated” form of traditional trial and error experimentation in terms of attempts at successful integration of both the complexity of a targeted intervention and its respective systems effects. If all goes well, they should lead to beneficial clinical outcomes in combination, and this is what personalized approach is supposed to ensure: that the holistic data about the patient is taken into account when developing particular targeted interventions. This approach stems from systems biology and is largely reliant on genetic screening and high-throughput methods in order to find interventive targets, but the dosage and the variants of these

interventions will be highly dependent on the individual setup of the patient. The rejuvenated trial and error experimentation can be said to refer to a hit and miss strategy in integrating the complexity of the intervention with the system complexity of the individual patient. It also appropriates the first meaning, “trial and error” in clinical trials, since it is through adaptive design that the system response is accounted for.

I agree that trial and error experimentation is indeed a part of translational efforts, but mostly as ‘business as usual’ of basic preclinical and early clinical research, now incentivized by what van der Laan and Boenink (2015) call external measures: closer collaboration of researchers, better exchange of data and knowledge, and physical proximity of laboratories and clinics. Insofar as translational medicine is about trial and error experimentation and understanding mechanisms, it is nothing scientifically new, but it is a structured effort to foster serendipitous discoveries by a better exchange of people, knowledge, and material. Finally, trial and error method is rather time-consuming, unless enhanced by more resources. In this respect, translational medicine is offering the necessary infrastructure to achieve the envisaged goal of medical innovation. Enhanced logistics, better coordination of different stakeholders, and more communication between different disciplines is very much the external, contextual factor that is called on to make a change.

However, contemporary biomedicine has at its disposal a largely expanded material base of science which consists of data banks and new technologies for analyzing them. For example, biobanks provide a more reliable access to the molecular base of diseases than animal studies. In addition, they invite new methods for dealing with such a large amount of data. High-throughput screening technology can find appropriate targets for intervention by examining large numbers of targets in parallel. It cannot, however, predict systems effects after the intervention has reached the targeted molecular base.

“The systems biology approaches, by contrast, seek to include rather than exclude the complexity of the disease biology and count on unexpected “hits”. Since it is taken to be highly problematic to infer biological effects from molecular action and since important “emergent properties” are seen to arise on the systems level, the study of isolated targets is considered to be largely futile. The systems biology approaches are directed towards drugs that act on many targets or towards combinations of specific drugs with the potential for “more-than additive” effects. Altogether, the novel

approaches aim at controlling the complexity of biological systems through complex interventions.” (Adam 2011, 78)

Mathias Adam (2011) argues that successful and efficient drug development demands that the two levels of complexity are accounted for in the development process from the start (p. 67). Due to their fundamentally different principles, “drug development is captured in a methodological dilemma that is considerably obstinate” (p. 68). He calls the two different kinds of complexities “complexity of (local) intervention” and “complexity of systems effects” (system complexity). Examples of local intervention complexity are high throughput screening technologies and rational drug design (the blockbuster drug discovery method of 1970s and 1980s). Examples of systems complexity are traditional empirical strategies (trial and error experimentation) and systems biology approaches, prominent in contemporary personalized medicine.

“In general, since systems effects remain largely unpredictable from local interventions, the modeling and testing of drug-target interactions cannot “reach up” to systems effects, while the investigation of systems effects cannot be tracked down to a molecular level on which it could direct the chemical design of drugs. This dilemma sets limits on the degree to which drug discovery and development can be turned into a systematic enterprise at all, for instance by being guided by a scientific understanding of underlying mechanisms or by exploring the options for pharmaceutical intervention in a methodical way.” (Adam 2011, 68-69)

Personalized approaches include the “rejuvenated” trial and error experimentation which integrates target searches with unexpected hits on the systems level. They include internal factors that accelerate discovery, make research more translatable, and therefore serve translational aims very directly. First, on the level of basic research, it is the omics advances and systems biology. Second, on the level of clinical research, it is the emergence of adaptive trials. However, I have argued that these scientific transitions are strongly facilitated by contextual factors which are not inherent to science itself, but entangled with it.

Translational medicine boils down to either external measures which offer nothing genuinely new, or to something genuinely new which is nowadays largely covered by the personalized medicine approach. These novelties, however, are not a direct scientific consequence of translational medicine, but rather of advances in basic science, systems biology, and technology development which were institutionally brought together under the umbrella term of

translational medicine, in order to provide the infrastructure that would purchase on the advances. In doing that, the line between what is internal to science and what is external has been erased. Important scientific advances can only become important drivers of societal change if they are supported enough to come to prominence. These “synergies” and “catalytic” nature of “unique” and “interdisciplinary” projects (NIH 2014) supported by government money and private-public partnerships lie at the heart of translational medicine. Personalized medicine, in particular, is dependent on highly optimized chances for successful hits and integration of different biological and environmental levels. Chances for success are optimized but still dependent on serendipity, similar to the paradigmatic case of penicillin, but largely enhanced as compared to traditional trial and error.

According to an account of serendipity put forward by Samantha Copeland, serendipitous discoveries include chance, sagacity, and a valued outcome (Copeland 2015, 28-47). She argues for an understanding of serendipity in science as

“...an emergent property of scientific discovery, describing an oblique relationship between the outcome of a discovery process and the intentions that drove it forward. The recognition of serendipity is correlated with an acknowledgment of the limits of expectations about potential sources of knowledge.” (Copeland 2019, 1)

The general idea is that serendipity is retrospectively characterized as an unexpended insight, but it can be cultivated by epistemic communities. Moreover, there has to be an element of sagacity, a wisdom in recognizing that a new piece of evidence or a hypothesized connection should not be ignored. According to Copeland, serendipity in science has further features: it comes in variations so, for example, some serendipitous discoveries are single events, like the penicillin discovery, while some are processes taking place in communities, as it is the case in contemporary research networks. Serendipitous discoveries are contingent on contextual factors and inherently unpredictable. (Copeland 2019, 29). There is also a surprise factor in serendipity:

“Serendipity occurs when the limitations of epistemic expectations are exposed: a discovery is serendipitous because it arises from an unexpected source of knowledge, or because knowledge is produced in an unexpected way.” (ibid., p. 2)

Related to the requirement of the unexpected source of knowledge, Copeland argues that intermediate and early results of experiments and trials should be made available so that others can take them up and possibly use them as evidence for their claims, in this way fostering

serendipitous discoveries. This is very much in accordance with the variety of evidence that plays a role in personalized medicine approaches, and with adaptive design trials in which certain parameters can be changed in order to make them more fitting to the needs of patients. Also, targeted interventive practices in personalized medicine are more predictive than the ones in the usual drug discovery process. This is because in assessing the system response and deciding on best interventions for a particular patient, everything is taken into account, from patient's health history to her eating habits, genetic setup and reported work-life balance. Unexpected evidence of correlations is both a norm and is fostered in a way which is in accordance with Copeland's definition of serendipity in science. The unpredictability of drug discovery process in terms of integrating interventive complexity and systems complexity is exactly the kind of limitation of epistemic expectations that can most efficiently be overcome by cultivating serendipity in a scientific community. Drug discovery acceleration in translational medicine is best described by systemized efforts to cultivate serendipitous discoveries by integrated external and internal measures.

6.2 Measuring the effectiveness of medical interventions

On the level of clinical research, measuring the effectiveness of medical interventions is a highly time-sensitive practice, since not every condition has the same time span of development, remission, and possible relapse. The same holds for testing for safety and efficacy, and monitoring for side effects in drugs. Time-insensitive practices can contribute to an overestimation of the effectiveness of medical interventions, as well as to an underestimation of harm. Stegenga (2015b) gives an example of testing high-dose chemotherapy for breast cancer 18 months after the treatment.

“This temporal range was adopted from blood cancers, in which high-dose chemotherapy is effective. After 18 months it appeared that the high-dose chemotherapy had prevented recurrence of breast cancers. However, breast cancers grow slower than blood cancers, and so 18 months was an inappropriately short time to measure the outcome of the therapy. Later studies that used a longer temporal range found that high-dose chemotherapy did more harm than good for breast cancers. The physiological difference in growth rates between cancer types explains why high-dose chemotherapy is more effective in blood cancers than in breast cancers: since chemotherapeutic drugs operate by interfering with mechanisms of cell division, cells that divide rapidly are

more susceptible to chemotherapy (slower growing tumors are less susceptible to chemotherapy). The initial studies that suggested that high-dose chemotherapy is effective for breast cancer employed an instrument that was not sufficiently sensitive to the temporality of the disease.” (Stegenga 2015b, 64)

This case is illustrative of the plausible view that effects cannot be adequately assessed in a short time. However, it is not easy to establish an appropriate time frame for a research practice, since cases vary significantly. Nonetheless, “both literature and policy tend to assume that speedy translation of research into practice is a good thing. Delays are seen as a waste of scarce resources and a sacrifice of potential patient benefit.” (Morris et al. 2011, 510)

I have already stated the tension related to hindsight knowledge: the fact that we want speed when it brings about the better, such as in the case of Ebola vaccination trial, and that we do not want it when it brings about the worse, such as in Ritalin trials for children with ADHD or in high dosage chemotherapy trials for breast cancer. Stegenga (2015b) draws particular attention to three epistemological problems with measuring the effectiveness of medical interventions: the selection of a good measuring instrument, the use of an appropriate outcome measure, and the employment of a reliable method of extrapolating measurements from an experimental setting to a broader setting.

“The way these challenges are met in contemporary clinical research is unsatisfactory, which systematically contributes to overestimating the effectiveness of medical interventions.” (Stegenga 2015b, 71)

I focus more generally on the idea that research should and could be accelerated. Morris et al. (2011) conducted a review of literature describing and quantifying time lags in the medical research translation process. They were motivated by the fact that several policy documents were using a supposedly replicated estimation that it takes on average 17 years for research to reach clinical practice.

“Such convergence around an ‘average’ time lag of 17 years hides complexities that are relevant to policy and practice which would benefit from greater understanding” (Morris et al. 2011, 510-511)

Their aim was to compare empirical data on translational lags with the conceptual model (from T1 to T2’) of translational research. This was done in order to provide an overview of estimated time lags and where they occur (ibid.). The authors of the study started from the fact that each phase of translational research is associated with a lag (i.e. each bridging of a translational gap,

as discussed in section 2.1), and the fact that lags are either inevitable because they are necessary for good science, or they are a result of “non-value adding waits” (ibid., p. 511). They found that it was very hard to assess translational lags comprehensively since the literature on gaps uses proxy measures. For example, it focuses on dissemination and publication in peer-review journals as these are the most measurable ways of estimating lags. However, if there are significant lags in raising funds or getting ethical approval, it is often not reflected in total lag estimations. Because of the variations in proxy measures, it turns out that studies are never measuring the same thing, and therefore the estimated 17 years becomes a problem, especially since it informs policy measures that aim to lower the average.

“There also appears to be a mismatch between conceptual models of the translation process, and the measuring of lags. For example, the gap between guideline publication and translation into actual practice is often ignored, suggesting an under-estimation of the time lags in some cases. On the other hand, interventions may come into use before guidelines outlining them have been published – suggesting an overestimation of time lags in other cases.” (Morris et al. 2011, 518)

The conclusion of Morris et al. study is that the complexity of translational lag measurement is not adequately represented because of the preference for ‘averages’, with no understanding of distributions and variations. Since some lags are indeed necessary to ensure the safety and efficacy of therapies, inadequate understanding of where exactly the non-value adding gaps are, “‘blindfolds’ investment decisions and risks wasting efforts to reduce lags” (ibid.).

“The discussion in the literature fails to consider what is necessary or desirable, tending to assume that all lags are unwelcome. A key question for policy is to identify which lags are beneficial and which are unnecessary, but to answer this question it is necessary to have an accurate and comparable estimate of the lags.” (ibid.)

The Cooksey Report, for example, uses 12 years as the average time of translation of a therapy into practice, which is taken from DiMasi et al. (2003) (Cooksey 2006, 106) and depicts the time it takes to develop a new molecular entity (DiMasi et al. 2003, 164, 181). Another study, which estimated the economic benefit of cardiovascular disease research in the UK, used the estimate of 17 years to calculate the return based on data between 1975 and 2005, and found an internal rate of return being only 39 percent, which was considered unsatisfactory (Morris et al. 2011, 510.) It is clear that policy decisions should be based on best evidence and that current variations in lag measurement and lag identification do not provide reliable guidance. This

makes the desired goal of 5-10 years to implementation, as stated by the NIH Roadmap, highly problematic.

However, 5-10 years goal does not seem to be so unrealistic anymore, which is largely due to the transition towards personalized medicine that does not involve the same clinical path as traditional medicine. These changes need to be highly scrutinized, and a better understanding of the reasons for time-lags, as well as a better assessment of therapy effectiveness is necessary if biomedical research and development is to be conducted in an epistemologically and ethically sound way. Better predictive outcomes based on integrated data and combined targeted and systems approaches can help in dealing with inappropriate temporal ranges for assessing efficacy, but only to a limited degree, since the trade-off will remain: stopping a trial too early, hence risking the cost related to under-evaluation, or stopping a trial too late, hence risking the cost related to over-evaluation, i.e. unnecessarily prolonging the trial.

7. Ethical considerations related to translational medicine

I now take for granted that adaptive trials developed as a part of personalized medicine approach to therapies are a prime example of translational medicine, taken the goal of acceleration as the primary criterion. To claim that they are an improvement on RCTs would be very speculative in the absence of a reliable empirical assessment of the outcomes of both kinds of trials. The usage of adaptive design has been considered justified in cases of urgency which require quick action, like the Ebola epidemic from the introduction of the thesis. Adaptive study showed efficacy of the experimental vaccine and the therapy has been approved, but not by a full regulatory process and only for “compassionate use” in outbreaks. Furthermore, adaptive design trials are sometimes necessary in clinical research of treatments for rare diseases because it is not possible to obtain big samples of experimental subjects.

However, transition towards adaptive studies more generally might be opening a dangerous leeway for biases that were attempted to be blocked by blinded randomization inherent to RCTs. A response to this concern is that genetic interventions have more predictable efficacy since they target specific molecular mechanisms and integrate a network of specific individual information in assessing the dose and type of intervention. Nonetheless, biases connected to commercial interests are not avoided by the fact that companies do not have to conduct huge and long trials, since the cost of personalized therapies is largely increased compared to chemically synthesized drugs that have been the norm not long ago. The interest in pulling the

treatments through is the same. Moreover, preventive practices like genetic testing pose significant ethical problems connected to data privacy, risk estimation and risk communication. Finally, taken the current socio-economic context, personalization means that accessible implementation in T2 is basically lost from sight as an achievable goal, since personalized medicine is not available to most people, though it seems to do a great job in T1 translations as far as we can access their effectiveness. Ethical discussion of various aspects of personalized medicine and adaptive trials therefore needs to be continuously encouraged and undertaken alongside the technological advances in an inclusive and multi-perspectival way. I would like to recall Maienschein et al.:

“The problem is taking the translation as an unquestioned desirable goal and trying to make the ethics fit.” (2008, 50)

When it comes to the ethics of translational medicine, there are several contributions in the literature addressing various aspects of it, such as Kagarise and Sheldon (2000), Maienschein et al. (2008), Sofaer and Eyal (2010), and Kimmelman and London (2011). I have already shared the worries expressed in Maienschein et al. (2008) related to translational ethos that takes application as the goal at the possible expense of distorting valid epistemological means to achieve it. A further concern of Maienschein et al. (2008) is that the unquestioned goal of application will accelerate medical innovation while at the same time obstructing the ethical discussion. Their criticism is grounded on the analysis of the NIH Roadmap objectives, as well as on stem cell research as a case study in translational medicine.

We could claim that RCTs are the most valid method of evaluating new therapies, while the acceleration goal promotes adaptive measures which can be seen as “distorted epistemological means”. On the other hand, such distortion is justified under one or both conditions: of urgency (non-epistemic), and of better understanding of the intervention and hence improved prediction of its effects (epistemic). However, acting under either of the two conditions (or both) will still involve trade-offs: between understanding and intervening, and reliability and speed. Excluding biases and confounding factors is an important indicator of reliability and leads to better understanding of the intervention. RCTs will promote this aim very directly. When the goal is to intervene, like in the Ebola case, randomization will be biased in favor of an increased access to the experimental treatment, which might be done at a cost of understanding. We have seen how AZT was shown to be harmful in the longer follow up trial. On the other hand, personalized therapies which make use of adaptive design do not aim at highly generalizable findings, since

their primary asset is a focus on the particular, rather than the general. Moreover, their focus is very clearly on intervention.

The general impression from scholarly work on the ethics of translational medicine which focuses on translational practice is that there is a confusion about what translational medicine is. Not surprisingly, since it is indeed an elusive phenomenon. For Solomon, it is application-driven basic science, for Maienschein et al. (2008), Fang and Casadeval (2010) and Jogalekar (2012) it is applied research without adequate fundamental knowledge acquired in basic research. Finally, translational medicine can also be captured by Berwick's "pragmatic science" where it amounts to clinical research. Most of the time, it is the uninformative "bench to bedside (and back)". I have settled for the criterion of acceleration and argued for personalized medicine as representative of translational efforts. If I am right, that only speaks in favor of the overarching scope of the label.

Mary Jane Kagarise and George Sheldon (2000), for example, identify translational research with clinical research in their paper "Translational ethics: a perspective for the new millennium" According to them, translational ethics is "based on autonomy and informed consent" which "helps navigate the ethical ramifications of technological and scientific advances that will increasingly challenge the corporate-oriented health system in the new millennium" (2000, 39). Their concern is that the rise of emerging technologies whose risks and benefits are not well understood will impede informed communication and autonomous decision making, so they highlight the important role of these concepts in clinical settings.

Neema Sofaer and Nir Eyal (2010) take a different approach and praise translational medicine for its attempts at implementation in T2. While agreeing with concerns about "questionable benefits, special risks, additional barriers to informed consent, and severe conflicts of interest" (2010, 19), Sofaer and Eyal shift focus to the prospects of translational medicine to address ethical problems related to research conducted on "global poor", and at the same time their limited access to the products of research. They recognize the difference between the two translational gaps and agree with concerns around research ethics in T1, which amount to the impossibility of predicting and communicating harms in translating from pre-clinical studies to studies in humans. However, when it comes to the renewed attention to T2 in translational approaches, such as improving health services, prevention, and implementation, they argue that translational efforts should be supported.

“Therefore, insofar as T1 is morally problematic, and no independent objections to T2 exist, the ethics of translational research is diverse: while some translational research is problematic, some is not. Funding and oversight should reflect this diversity, and T2 should be encouraged, particularly when conducted among the global poor.” (ibid.)

The authors are calling for a better implementation of therapies, and their argument is made from a global health care perspective which identifies the strengths of translational efforts on the T2 side. We have, however, seen from both van der Laan and Boenink (2015) and Solomon (2015) that translation is mostly identified with T1. Mentions of T2 are rarer and translational efforts are apparently less represented in T2. Personalized medicine is especially unlikely to become available to everybody any time soon.⁶⁵

Kimmelman and London’s paper (2011) “Predicting Harms and Benefits in Translational Trials: Ethics, Evidence, and Uncertainty” focuses on early clinical trials and the problem of making their outcomes more predictable. This is a discussion related to the role of mechanistic evidence in trials, as well as to ethical questions regarding human experimentation under uncertainty. The authors argue for increased reporting of animal studies outcomes in order to predict the outcomes of early human studies better, since they present a risk for the participants. In Kimmelman and London’s paper, ‘translational trials’ are not characterized by any adaptive measure. They are translational because they are first in-human trials after animal studies, i.e. they are T1 translations. The fact that these trials are ‘translational’ does not mean that they were inherently different before translational medicine came to prominence, it only means that the term has now gotten into widespread use. When translational era had been on its peak, the researchers were likely to use the label. The stem cell researcher in Harrington and Hauskeller (2014) was explicit in linking the term with better prospects for getting funding.

Finally, an especially insightful contribution at the intersection of ethical and epistemological considerations comes from Lieke van der Scheer et al. (2017) in a paper called “The Benefits

⁶⁵ I have recently been able to follow a huge fund-raising campaign in Croatia to recover the cost of treatment for two-year-old Mila who is suffering from a rare type of leukemia for which a personalized therapy is available in a clinic in Philadelphia at a cost of 2.8 million dollars. More than 5.8 million dollars were raised and Mila is currently responding well to treatment in America, while the rest of the money is reallocated to a fund for children in similar situations. The money was collected completely through citizen action, since no institutional means at disposal could regularly provide such amounts for individual cases (see All for Mila, n.d.).

of Patient Involvement for Translational Research”. The authors are also at a loss in terms of what translational research precisely is, since one of the prerequisites for the arguments they make is “to establish the main perception of TR” (ibid., p. 3). The paper is motivated by the insight that “patients, who are the original object of concern in biomedical research, do not play an active role in the discourse of translational research” (ibid.). This is epistemologically problematic, since patient contribution can improve science, but it is also ethically problematic, since patients are the most vulnerable party in the network of stakeholders in biomedical knowledge production, and their voices should be encouraged and heard in the supposedly interactive exchange that is promoted. The question that is raised is “whether patient involvement would be a beneficial way to help determine and achieve the aims of TR and, if so, how to proceed” (ibid.). This is a very interesting aim 13 years after the Roadmap, which has been set to revolutionize the way biomedical research is done and in which better patient outcomes were one of the goals. I have already argued that bringing a drug to the market does not necessarily correlate with better patient outcomes or improved healthcare.

Van der Scheer et al. (2017) give two main arguments for patient involvement in translational medicine. The first argument has to do with empowerment, democratic rights, and legitimacy. Patients should have an active, rather than a passive role in bringing about the change, i.e. they should be empowered as participants and negotiators in a research process which deals with their interest. This also comes from their democratic right to participation. Patient participation enhances the legitimacy of political decisions regarding which research directions to pursue, because patients are “the ‘owners’ of publicly funded research” (ibid., p. 8). The second argument for patient involvement comes from their “situated”, experiential knowledge, which can contribute to the research process (ibid., p. 10).

They give two examples of successful integration of patients as partners in research. The first is the already mentioned AIDS activism, where the research process was influenced by patients’ experiences and involvement in questions such as which research direction should be funded and how to define AIDS. The definition of AIDS was changed in order to incorporate HIV-related conditions that affect women, as a result of patients’ participation (ibid., 11). Another example is from a conference on Outcome Measurement in Rheumatology Clinical Trials (OMERACT) where patients succeeded to put the issues of fatigue, general well-being, and sleeping problems on the research agenda and bring about a change in outcome measures (ibid., p. 11).

Contrary to that, there is also evidence of unsuccessful attempts to integrate patient perspectives. One of the perceived problems is the already mentioned overemphasis of RCTs in evaluating therapies, which often do not adequately incorporate patients' values. Furthermore, patients' voices often get excluded due to a lack of formal expertise, even when the environment for dialogue is supposedly inclusive. In other words, patients are called to the table, but their voices are not properly heard because they are outnumbered or not listened to by the doctors and researchers who are often not likely to actively involve patients as equal contributors to research. Bueter (2019) argues that the exclusion of patients as partners in research practices can be considered a form of epistemic injustice, namely pre-emptive testimonial injustice. It means that patients' opportunity for testimony is denied due to its wrongly presumed irrelevance or their lack of expertise. Van der Scheer et al. emphasize that in those cases where integration was successful, it was because "the patient group was well defined, well organized, and the patients were competent and had a knowledge of and were well trained in research methods" (ibid., p. 11).

Finally, van der Scheer et al. identify an important obstacle for a successful integration of patient experience in translational research, and that is an obstacle grounded in normative claims of translational discourse:

"However, the question remains whether patient involvement and the contribution of experiential knowledge of patients are always compatible with the normative goals of translational research, i.e., to more quickly improve a patient's quality of life in an economically sound way. If it were compatible, one would benefit from this nice win-win situation. However, the fact remains that, though it may occasionally be fortunately true, this is not always the case." (ibid., p. 10)

Patients' involvement and inclusive negotiation takes time and effort, which is the most direct obstacle taken the aims of translational research. However, the authors' argument is not pessimistic. Taken the benefits of patient involvement, van der Scheer et al. believe that including them early enough would help to shape research more directly towards patients' needs, and thus contribute to the goal of acceleration, relevance, and usefulness.

"The discourse concerning TR is inherently normative because it is motivated by the number of values it wants to achieve such as relevance, usefulness, and the economic value of the innovations. Whereas the standard view of TR suggests that these values may be achieved by facilitating the step from lab and animal studies to clinical trials by

removing obstacles external to science, alternative views of TR point out that the innovation process is multidirectional, also includes the steps from knowledge and technology to employ in practice, and pays attention to obstacles both outside and inside scientific practice. Beginning from this more complex view of what is required to further translation, it appears plausible that early involvement of future users, in particular patients, during the innovation process could have a positive impact on the achievement of those values.” (ibid., p. 13)

Though it takes time to include patients, their involvement might be reducing the overall time of therapy development, since it would make products more fitting to their needs. This suggestion aligns well with my view that personalized medicine is the natural successor of translational efforts. One of the “Ps” in “P4” motto of personalized medicine is “participatory”. Personalized medicine is centered on taking patients’ needs seriously, at least when it comes to their individual therapeutic needs. It yet has to be assessed to which extent the values of patients are meaningfully included in the process of research, and how the risks and benefits are communicated. It is outside the scope of this thesis to give an assessment of participatory practices in personalized medicine. One of the reasons why it is hard to make an assessment is because no long term effects of recently approved personalized medications and practices (that can be said to be accelerated as a consequence of translational efforts) are available as evidence. This being the case, it is clear that efficient communicating of risks and benefits poses serious problems for ethics, as well as for regulation. Furthermore, personalized therapies do not only amount to life saving drugs, but also to a myriad of preventive genetic tests, whose benefits have to be carefully scrutinized. We have seen how incorporating patients’ values can be seen as contrary to the goal of acceleration, but also which benefits are gained by increased patient involvement. It is thus necessary to pursue efforts towards better patient inclusion in the research process.

8. Research on cortisone in the 1930s-1950s⁶⁶

Before concluding the case of translational medicine, I would like to present a less-discussed case of a paradigmatic translation in the pre-translational era, in order to make vivid several

⁶⁶ A version of this chapter is published as a review article in *Acta medico-historica Adriatica* (Jadreškić 2016).

characteristic features of biomedical medicine as it used to be done. Historical success stories have inspired contemporary translational discourse, but the context in which medicine is done today is radically different. However, despite differences, some lessons pertaining to problems with acceleration can still be learnt. I will use a historical case to underlie some common problems of accelerated translations, and at the same time provide an insight into how biomedical research was done only several decades ago, which might be helpful for putting contemporary practices in a larger socio-economic context. The larger context will then be taken up in the conclusion, which will link contemporary biomedical research with more general problems of values in science. To values I will attend in the third part of the thesis.

8.1 Introduction

The discovery, synthesis and therapeutic application of cortisone present a paradigm for modern translational medicine according to Hillier (2007) and Saenger (2010) because they represented a joint achievement of discoveries by biochemists, Edward Calvin Kendall and Tadeus Reichstein; large scale synthesis by an industrial chemist, Lewis Hastings Sarett, and therapeutic application by a rheumatologist, Philip Showalter Hench. I will present conditions that made this basic/applied/clinical research interface possible: the rise of steroid chemistry, simultaneous individual accomplishments as well as continuous cooperation between scientists, military competitiveness, and cooperation between pharmaceutical companies.

Since their first medical use in 1948 cortisone and its synthetic analogues have remained among the most widely prescribed medications in the world (Hillier 2007). Cortisone belongs to the group of steroid hormones of the adrenal cortex, first isolated in the 1930s by biochemist Edward Kendall in the USA and independently by Tadeus Reichstein in Switzerland at about the same time. The discovery of cortisone turned out to have a medical, scientific and industrial importance, and it led to further discoveries with wider implications for drug development, such as conformational analysis important for the later emergence of rational drug design (see Quirke 2005; Slater 2000).⁶⁷

⁶⁷ Conformational analysis is a method for correlating steroid structures with the physical and chemical properties of the molecules (see Quirke 2005, 647; Slater 2000, 471-480).

Steroids are a group of compounds with a common structure based on the steroid nucleus, consisting of three six-membered carbon rings and one five-membered carbon ring, which occur in plants and animals (Slater 2000, 444). The general name steroid was introduced in 1936 to cover all compounds with a steroid-like skeleton. The sex steroids estradiol, testosterone and progesterone were discovered between 1929 and 1935. What followed is the discovery of adrenocortical hormones between 1935 and 1938 (Hillier 2007, 1). They were produced by partial synthesis, i.e. synthesis that began with structurally elaborate starting materials, complex natural products. A total synthesis referred to a synthesis starting from simple and inexpensive materials whose composition was known, such as air and coal. Robert Burns Woodward achieved a total synthesis of cortisone in 1951 (Woodward, Sondheimer, and Taub 1951). The discovery of cortisone and its therapeutic efficacy led to what Hillier called the “diamond decades” of steroid chemistry that started in the 1950s (2007, 2). Cortisone success, however, is often described as a mixture of knowledge and luck:

“A fascinating tale of good science, perseverance, and luck which might not be possible in today’s regulatory environment.” (Burns 2016, 1)

Leo Slater (2000) used cortisone research to show that the boundaries between the disciplines were flexible and how the networks of research interpenetrated one another. I will mostly follow his work on the history of cortisone.

8.2 Discovery and translation

Edward Kendall at the Mayo Clinic in Rochester, Minnesota and Tadeus Reichstein at the Eidgenössische Technische Hochschule (ETH) in Zürich independently identified the structure of cortisone in the 1930s from extracts of bovine adrenal glands (Mason, Meyers, and Kendall 1936; Reichstein 1936). Kendall isolated eight crystalline cortical compounds from about 1.250.000 cattle carcasses, while Reichstein isolated twenty-eight crystalline compounds from the adrenals of 20.000 head of cattle (Slater 2000, 451). Kendall had the assistance of commercial agreements and contracts with Parke-Davis in Detroit and Wilson Laboratories of the Wilson Packing Company of Chicago which provided some 150 tons of adrenal glands for his research between 1934 and 1949 (ibid.). Kendall also produced about \$9 million worth epinephrine, sold as Adrenalin, for Parke-Davis (ibid.).

“Kendall's program was a mixture of industrial process, chemical investigation, and basic medical research. Nothing was inherently pure or applied about this biochemical activity.” (ibid.)

At the same time Reichstein worked with crude extracts provided by Organon, Inc. of Oss, Holland, not with adrenals (ibid). Reichstein and Kendall independently assigned their new compounds letter designations; Reichstein's “substance Fa” and Kendall's “compound E” later proved to be identical, and Kendall named them cortisone in 1949 to avoid confusion with vitamin E (ibid., p. 451-452).

Independent biochemical investigations were still far from discovering therapeutic efficacy. This was the accomplishment of Kendall's colleague, Philip Hench, the head of the Department of Rheumatic Diseases at the Mayo Clinic. Hench hypothesized the presence of a therapeutic agent X that emerged in conditions of jaundice and pregnancy, and relieved arthritic patients from rheumatic symptoms (Hench 1964). His physiologic research of an unknown therapeutic agent began in 1929, and it took twenty years until his announcement that certain corticosteroids were able to reverse many of the acute manifestations of rheumatoid arthritis (Quirke 2005, 646). Simultaneous discoveries made the identification of cortisone characteristics possible. One was biochemical, related to the steroid structure of the hormones isolated from the adrenal cortex, and the other was physiologic, related to therapeutic possibilities for rheumatic patients. The discovery of a biological therapeutic substance was undertaken in ignorance of its site of origin, relying solely on its function in the organism. Hench eliminated the possibility that agent X is a sex hormone, because patients of both sexes were relieved from arthritis pain with the occurrence of jaundice. Successful converging of the two discoveries made at the same clinic presents the special serendipitous moment in the cortisone chronology.

The outset of World War II played a significant role in the reinforcement of collaboration between scientists on the issue of adrenocortical hormones. In 1941, as the United States became involved in World War II, adrenal research became an “internationally competitive effort” (Simoni, Hill, and Vaughan 2002, 21). As Viviane Quirke notes:

“The therapeutic potential of adrenal cortical hormones in rheumatic and other inflammatory diseases might never have been investigated, had it not been for the rumour in 1941 that Luftwaffe pilots were taking these hormones to increase their resistance to oxygen deprivation and be able to fly at higher altitudes.” (Quirke 2005, 649)

A German submarine was captured with cargo that was believed to be adrenal glands. A report saying that the Germans were buying adrenal glands in Argentina reached Washington (Slater 2000, 452). It was later revealed that the cargo in the submarine was in fact liver for Otto Bayer's work on vitamin B12 at IG Farbenindustrie (Quirke 2005, 649, footnote 29).

It is true that steroids were being produced out of extracts of sex organs between 1920 and 1940 in Germany, but none of them was cortisone. The first European adrenal product was Doca (Deoxycorticosterone Acetate), a chemical sold by Swiss Ciba, identical to the later Cortiron, launched by the German Schering (Gaudillière 2013, 191). Ciba became involved very early with steroid chemistry through its collaboration with the ETH in Zurich and later on with the University of Basel (Heusler and Kalvoda 1996). Ciba had an agreement with ETH in Zürich, and therefore benefited from Reichstein's work. All patent rights on steroids were assigned to Ciba. However, it was Schering that dominated the German drug market, and Dutch Organon which supplied Reichstein with research material (Gaudillière 2013, 191). So the three companies, Schering, Ciba, and Organon, signed an agreement in 1939 called "Cortin", which organized and limited competition among these corporations, providing important resources for the flow of research information and for licensing of patents (*ibid.*). Early uses of Schering's Cortiron remained limited because its usefulness was not clear, and cases of Addison's disease, that it was hoped to alleviate, were rare (*ibid.*).⁶⁸ That did not change during the war, even though steroids were given military priority because of their use in healing wounds (*ibid.*). We can see how the boundaries between basic and applied were vague, and how European pharmaceutical and academic communities exchanged information and material, and facilitated intensive research. The line connecting ETH, Organon, Ciba, and Schering can be traced to the University of Basel in Switzerland, where Reichstein continued his work.

In America, a similar cooperation of industry and academy took place. The rumour about German usage of steroids motivated the American National Research Council to rank steroids at the top of wartime research agenda, above penicillin and antimalarials (Slater 2000, 452). Merck and the Mayo Clinic, as well as other academic and industrial groups, were collaborators in the steroid research program set up by the Office of Scientific Research and Development (OSRD), which intensified the research on adrenals and promoted the crossings of disciplinary boundaries (*ibid.*). It was also in 1941 when the decision was brought to administer compound

⁶⁸ Addison's disease is a chronic condition of adrenal insufficiency, deadly if not treated with hormone replacement therapy.

E (cortisone) to a patient. Hench and Kendall met at a conference, where Kendall remarked that his compound E increased the resistance of animals against reactions to typhoid vaccine. Hench made note of this in his pocket notebook, and eight years would elapse before there would be enough of this substance to administer to a rheumatoid patient (Hench 1964, 318-319).

The priority of adrenal research was the reason why Lewis Sarett from Merck and Company worked for three months in Kendall's laboratory in 1942 (Slater 2000, 452-453). The idea was to gain knowledge and return to Merck with the goal of developing large-scale synthetic methods for the compounds that were chosen for the initial studies, compounds A and E, because of their relative structural simplicity. After several years funds were cut, as the effectiveness of the research for various war-related uses did not meet the expectations. Eventually, only Kendall's group at Mayo and his collaborators at Merck continued. By the end of the war, in 1945, enough compound A had been available for clinical testing. When it proved to be ineffective in the cases of Addison's disease, it was "a big disappointment" (Kendall 1964, 273).

Kendall and his associates, however, pursued the research. Sarett synthesized cortisone in 1946 in 37 steps at Merck (Sarett 1946). Shortly after the war ended Merck management decided to make cortisone available for clinical investigations. In September 1948 Philip Hench administered 100 milligrams of cortisone intramuscularly to a patient suffering from rheumatoid arthritis. It showed rapid success in relieving pain and reducing inflammation. Subsequent trials with cortisone and adrenocorticotrophic hormone (ACTH) gave similar symptomatic relief (Hench et al. 1949; Hench 1964).

The patent was assigned to the non-profit Research Corporation (Kendall 1964, Slater 2000). The Research Corporation had been founded in 1912 partly to prevent the commercial exploitation of academic research (Slater 2000, 457). The Corporation held patents for scientists at universities and foundations and licensed these patents, producing income for the support of other research, and providing for the redistribution of research income (ibid., p. 457-458). In August 1949 the Committee of the National Academy of Sciences on the Investigation of Cortisone posed the problem of distributing the limited amount of cortisone ("News and Notes", *Science*, August 1949). It was decided that the small amount available will be used only for clinical and experimental research and that it will be made available to selected investigators from the USA and Canada. Since the Academy had no funds with which to buy cortisone or to support research, it was:

“(…) confidently expected that the needed funds will become available from both public and private sources. (…) The Academy committee has accepted this responsibility because of the deep conviction that a new discovery of the greatest importance to the health and welfare of countless people has been made and that it is vital to promote its most rapid and intelligent development.” (ibid., p. 153-154)

In the autumn of 1949 the Research Corporation held the Cortisone Conference, where seventeen scientists to whom the Corporation provided support were invited (Slater 2000, 458). Topics included total synthesis, natural raw materials for partial synthesis, hormone analogues, toxicity, and clinical testing.

Kendall, Hench and Reichstein were awarded the 1950 Nobel Prize in Physiology or Medicine for their discoveries relating to the hormones of the adrenal cortex, their structure and biological effects. On November 1 1950, cortisone became available to physicians in the United States (Kendall 1964, 285). Merck's brand of cortisone was named Cortone. The price was falling: In July 1949, the price per gram was 200 dollars. In 1950, it was reduced five times, from 150 dollars in January to 35 dollars on November 1 (Kendall 1964, 276-277). However, in 1951 we can still read that cortisone “cannot yet, or perhaps ever, become a cheap medicament” (“Paths To Cortisone”, *BMJ* 1951, p. 408). This was mostly due to the lack of supplies of bile acid. Foreign sales were minimal, regulated by the US government, and were made primarily to those countries that would contribute the starting material, cattle bile, or participate in clinical development (Slater 2000, 467).

8.3 Implementation and problems

Jean-Paul Gaudillière called cortisone “an iconic product of the ‘therapeutic revolution’” (2013, 190). Cortisone was used in treating excessive inflammation, allergy, acute infections, and autoimmune disorders. Eventually, it had a deep impact on several medical specialties, including ophthalmology, gastroenterology, respiratory medicine, dermatology, nephrology, endocrinology and rheumatology. John H. Glyn, a British rheumatologist and Hench's friend and colleague, wrote in 1998 that the discovery of cortisone transformed rheumatology from its “Cinderella status of the BC (before cortisone) era” (Glyn 1998, 823), though its clinical usefulness remained controversial. Beneficial effects in rheumatic patients and its anti-inflammatory agency were challenged by serious side effects brought on by high dosage levels and prolonged therapy (Beckett and Stevenson 1956), as well as by relapse of the treated

condition once the treatment was ended (Wilson 1950; Carlisle 1950; and Thorn 1951). It was clear that “cortisone suppresses rheumatoid arthritis but does not cure it” (“Treatment With Cortisone”, *BMJ* 1951, p. 222). It was discussed with caution, considered to be “a real danger if (...) it is used anything like as empirically by medical men generally as were the sulphonamides and penicillin” (Emery 1949, 653). Glyn wrote in his memoir:

“In the United States a black market developed which had serious medical and social repercussions. Patients who had experienced great relief of their symptoms were not prepared to relapse when supplies ran out. They became totally dependent on the drug. Overdosage led to devastating side effects, and the ever escalating cost of maintaining their supplies resulted all too often in financial destitution. Such patients had no alternative but to seek relief by registering as guinea pigs to research groups such as the one at the Bellevue Hospital in New York which I joined in 1952.” (Glyn 1998, 823)

Alternative sources of starting material and new methods of production were investigated because of shortages of bile acids and of money (“New Sources Of Cortisone”, *BMJ* 1950). Kendall reported at the American Chemical Society meeting that 40 head of cattle were needed to provide cortisone needed daily by one patient (“In Science Fields”, *The Science News-Letter*, February 1950, p. 121). Merck officials expected the production to be tripled or quadrupled by the middle of 1952 (“In Science Fields”, *The Science News-Letter*, March 1951, p. 168). The increased output was expected from recently investigated botanical sources as well as from total synthesis.

John Glyn led the first UK cortisone studies. Small quantities of cortisone and ACTH were distributed to hospitals in Britain (“Cortisone And A.C.T.H. In Britain”, *BMJ* 1950, p. 1375). Hench gave two lectures in London in 1950 (*ibid.*) and Glyn published the preliminary results of the studies in the *British Medical Journal* in 1950 (Copeman et al. 1950), and the next in 1952 (Copeman et al. 1952). In 1954 the results of the first multi-center crossover trial of cortisone and aspirin organized by the British Nuffield Foundation and the Medical Research Council were published in the *British Medical Journal*. The trial concluded that there was no evident difference between the two groups. This led to a sharp correspondence between Glyn and Austin Bradford-Hill who designed the trial (Glyn and Todd 1954; Bradford-Hill 1954). Glyn criticized the design, the selection of patients, and the dosage:

“My comments on this trial are largely a text on which to hang a plea to the statisticians to modify their rigid approach to clinical trials. Perfect statistical techniques are not

possible when dealing with biological material, nor is it possible at the early stage at which the Medical Research Council generally formulate their trials for them to lay down the optimum regimes which are not liable to subsequent criticism. (...) In other words, the pendulum away from therapeutic empiricism has swung too far.” (Glyn and Todd 1954, 1376)

We learn that in the later years, in a personal conversation, Bradford-Hill “graciously agreed” that some of Glyn's comments “were justified in the light of subsequent events” (Glyn 1998, 823). Austin Bradford-Hill was an English epidemiologist and statistician who pioneered randomized control trials and established a criterion of causation (Bradford-Hill criteria of causation) in the pre-EBM era (see Howick, Glasziou, and Aronson 2009). RCTs have been the cornerstone of evidence-based medicine since then. Dissatisfaction with statistical techniques in Glyn’s 1954 letter to Bradford-Hill resonates with the contemporary criticism of EBM. Cortisone in particular has brought additional awareness of a need for individualized therapy and patient approach.

By 1955, analogues with reduced toxicity and enhanced physiological activity were developed (“Modifications Of Cortisone”, *BMJ* 1955, p. 1520). Although Kendall believed it “highly improbable” (Kendall 1964, 278) that any product would occur that would be used in the place of cortisone and cortisol,⁶⁹ generic formulations of prednisone, prednisolone (see Jenkins and Sampson 1967) and dexamethasone remain in widespread use to this day. In chemistry, a development of the method of conformational analysis facilitated the study of the relationship between structure and activity, enabling a rational approach to drug design (Quirke 2005, 647). In medicine and pharmacology cortisone was an important incentive for a new attention to chronic diseases that occurred in the post-war research (*ibid.*, p. 648). It is now known that intracrine⁷⁰ metabolism of cortisone to cortisol sustains local amplification of glucocorticoid action at sites of inflammation throughout the body (Hillier 2007, 1).

⁶⁹ Cortisol is the closely related compound F, in Kendall's notation.

⁷⁰“Intracrine” denotes a type of hormone function in which a regulatory factor acts within the cell that synthesizes it by binding to intracellular receptors (“intracrine.” 2007)

8.4 Discussion

Cortisone was discovered around 1936, while partial synthesis for large scale use was achieved in 1946, ten years later. First therapeutic application followed in 1948, commercial availability in 1949, and total synthesis in 1951. Taken these milestones in cortisone's translation to clinics, we get 12 years to first application (1948) and 15 years to a cheap method of mass production (total synthesis in 1951). This is below the widely adopted average of 17 years that Morris et al. (2011) have examined, and exactly as the average for new molecular entities adopted in The Cooksey Report. Still, Hench's work on an unknown therapeutic agent had started much earlier so it is again hard to estimate how long a particular research lasts and what can be considered as its starting points and milestones.⁷¹

The cortisone case is paradigmatic of translational medicine primarily because of the serendipitous discovery resulting from the cooperation between Kendall and Hench in the Mayo clinic. While Hench hypothesized an unknown therapeutic agent active in relieving rheumatic pain after the onset of jaundice and pregnancy, Kendall had a sufficient amount of steroid compounds ready for experimenting. It was crucial that they were able to meet, exchange information about their work, and set up a trial where these substances would be tested in humans. They first started with the simplest compound, since it was easier to get it in necessary quantities, and when it failed, they moved to the next simplest compound, which turned out to be successful. This was exactly the kind of trial and error experimentation that the usual drug discovery process was reliant on. There was not enough knowledge about pathophysiological mechanisms and they could only hope for a lucky guess, which was facilitated by an exchange of knowledge and material. In a dramatically enhanced form, I argue that a similar facilitation of exchange is happening in modern translational contexts as well. In the case of rheumatic pain, the only biomarker available was patient's reported condition, followed by the increased ability to move and eventually walk. The beneficial effects of cortisone were dramatic. The first treated person reportedly walked from the hospital just several days after the first injection.

However, what brought adrenal research into clinics in the first place was war-time prioritization of research efforts towards therapies that could help soldiers' endurance and pain relief. We have seen how the research on adrenals was turned into a competitive endeavor of highest national importance, and this is why Sarett from the Merck company was able to work with Kendall and learn about new advances in steroid research. Sarett's method of cortisone

⁷¹ Revisit the same argument in Part I, Ch. 2, p. 25-26.

synthesis pursued as a response to war-related incentives eventually provided Kendall with the necessary amounts available for experimenting, but only in 1946 when the war was over and the government interest in steroids had already waned. The acceleration of research efforts was facilitated by external measures (OSRD steroid research program) and because of social reasons, but in this case the driver of the research incentive was military usage and not the economic cost-effectiveness. Methods of production and related economic cost-effectiveness became an issue only after the discovery of dramatic therapeutic effects of cortisone. The research on steroids had again been incentivized, but this time the goal was to find methods of cheaper production. It was impossible to provide huge amounts of cattle needed to cover the requirements of production, especially with the increased demand.

This brought about an inversion of research efforts, in line with the interactive model of translational efforts discussed before. The search for knowledge, motivated by the desire to understand, has been the traditional realm of universities, while the practical aims of application and production have been considered inherent to industry. Steroids were synthesized by partial synthesis, starting from complex biological material whose availability was limited. A relatively efficient method of partial synthesis (37 steps) was developed in Merck with the aim to scale-up production. Although resources were limited, it was much easier to synthesize compounds from already complex materials, because the steps to final product were fewer than in total synthesis which is much more complex. Total synthesis had usually been pursued in academic contexts, because it takes more time, uses inexpensive materials, and provides more understanding, since it starts from the fundamental molecular structure of compounds. However, with the increased demand for cortisone, the research on total synthesis was now pursued in industry, because it would solve the problem of limited supply of cattle bile and make the method, though more laborious, incomparably cheaper. Total synthesis was eventually achieved by James Woodward in Harvard.

What I aim to show is how the understanding of usefulness is contextual on one's aims and resources, and how the relation between epistemic and pragmatic is interrelated. One could argue that partial synthesis is the more pragmatic way of getting cortisone since it is faster and less complex. However, the overall efficiency of this process depends on whether the resources are available. In this case it was a significantly less optimal method because it is dependent on the availability of animal material, unlike total synthesis. It was therefore pragmatic to pursue total synthesis, since it would solve the resource problem. Furthermore, if the epistemic aim is only to synthesize cortisone, both methods are good enough. But if the epistemic aim is to

understand the process of synthesis better, then total synthesis is preferred because it provides more insight into the details of molecular structure. If the material resources for both kinds of synthesis had been abundant, if everything else is equal, then the faster method, i.e. partial synthesis, would have been superior because it could produce more cortisone in a unit of time. But since they were not abundant, total synthesis had been the superior method because it allowed the usage of inexpensive material. Finally, if there had been a more significant difference in steps through which the synthesis is achieved between the two methods, it might have also changed the prioritization of one over the other. Compare this last case with the limitations in computation power that the gravitational wave researchers face. They of course prioritize what they can adequately analyze with the resources they have at a given time. These examples, again, support the thesis that epistemic aims are pragmatically informed.

The success of cortisone had another side though, which became especially apparent after it became widely available. The devastating and sometimes fatal effects of prolonged usage, as well as emerging cases of cortisone addiction, could not have been assessed immediately, since the effects of elevated cortisone in the body are cumulative. If cortisone is administered for a long time, it can lead to serious side effects, including glaucoma, fluid retention, high blood pressure, mood swings, and other psychological and neurological effects. It is also important not to abruptly stop a treatment with cortisone, since doses have to be gradually decreased. The enthusiasm of the early 1950s had thus been dropping as the subsequent years brought about problems of cortisone addiction and relapse of the treated conditions once the therapy was stopped.⁷²

⁷² The issue of cortisone addiction soon reached Hollywood. In 1956 a film called “Bigger Than Life” starring James Mason and Barbara Rush was released. The main protagonist is a school teacher and a *pater familias* who suffers from a rare condition accompanied by severe pain, which makes him start experimental treatment with cortisone. He becomes addicted and experiences severe mood swings and psychoses. The poster for the film illustrates a woman hugging a distressed man and above them a physician saying: “I prescribed it... he misused it!” (“Bigger Than Life”, n.d.) Such a description of events clearly suggests that the patient is blameworthy for the addiction, which is an ethically problematic interpretation. The question suggests itself: has this message been released through such an influential medium in order to influence public perception about what went wrong with the “wonder drug”? Such an interpretation might have been expected to alleviate the damage done by physicians’ insufficient knowledge about the effects of prolonged cortisone therapy. I do not have enough evidence to argue this, but it is not implausible.

Glucocorticosteroids, a group of steroids to which cortisone belongs, are now known to be an exception in dosing, because they are administered in high dosages initially, which are reduced onwards, while it is usually the opposite (Aronson forthcoming). The side effects of corticosteroid therapy are still a matter of dispute and controversy. In 2014 FDA issued a safety warning related to injections of corticosteroids into the epidural space of the spine, since they may result in “rare but serious adverse events”, including “loss of vision, stroke, paralysis, and death” (FDA 2014). Also, recent studies have shown that corticosteroid injections for knee arthritis decrease patients’ pain for a week, but have no effect in the longer run. This effect was overlooked because the conducted trials have not lasted sufficiently long to reveal it (Stegenga 2017, 26). Cortisone based medications are nonetheless an indispensable treatment in many diseases and conditions, in many cases lifesaving.

My aim was to show how a relatively (or agreeably) fast translation was achieved in the cortisone case and to which subsequent problems it led, despite of apparent success in treating patients. Several other features of cortisone translation are worth noting. We could see that the transgression of academia-industry boundaries had been usual, and that mutual exchange of material and information had been beneficial for research, but also limited by agreements, patent rights transmission, and by third party involvement (Research Corporation). We could also see how war time context prioritized research directions and how methods of cortisone production had been changing depending on the circumstances and the demand. In contemporary translational vocabulary, it was the partial synthesis that enabled T1, but it was the total synthesis that enabled T2. This is an example of an overlap between the internal and the external factors in research. The development of new methods is the internal, scientific factor of change. The demand, as well as financial or other incentives are a contextual factor. In this case both played a role. Total synthesis was pursued, among else, as a response to increased demand.

There are also factors inherent to the context in which science is done that differ from the contemporary context of translational research. The non-profit Research Corporation, for example, was the owner of patent rights, and although the academy-industry cooperation was usual, the commodification of knowledge was much less prominent. Furthermore, early experimentation was done without much regulatory supervision, which in this case turned out to be successful, but it could well be that early clinical trials not only ended without signs of efficacy, but also with significant harm. In this case, harmful side effects occurred much later than the first administration. Finally, the 1954 correspondence between John Glyn and Austin

Bradford Hill resonates particularly well with the contemporary debates about the limits of RCTs. Bradford-Hill was trying to establish a scientific statistical method for assessing effectiveness, which was later especially reclaimed in the evidence-based movement. On the other hand, Glyn calls for “therapeutic empiricism” which seems to be similar to what Solomon calls “trial and error in interventions”. Glyn especially relates therapeutic empiricism with dealing with biologic material.

This brings us to the origin of biomedical research movements, drug discovery methods, and clinical research design that I have discussed earlier. The research on cortisone led to one of the first total synthesis of steroids. Total syntheses of steroids enabled linking structural properties of molecules with their functional properties, which contributed to the development of the dominant drug discovery method in the later years – rational drug design. Nowadays we see a restoration of research on biological material, only now it is on the level of genetics and genomics, and through personalized approaches which reclaim trial and error experimentation. The rejuvenated trial and error experimentation optimizes chances for a successful integration of targeted interventions, characteristic of rational design and high throughput screening, with favorable system responses. We are also witnessing an emergence of adaptive measures in the design of clinical trials. Glyn’s statement that “perfect statistical techniques are not possible when dealing with biological material” (Glyn and Todd 1954, 1376) resonates well with adaptive measures, and against the gold standard RCTs of the recent decades. Personalized approach is especially indispensable in decreasing the risks of corticoid therapy by careful reduction in dosing.

The circle of biomedical epistemic practices is thus closed. Adaptive trials, personalized medicine, trial and error experimentation, but also statistical methods of RCTs, have their origins in earlier discovery and clinical practices. However, in addition to significant scientific advances in the course of the last decades and the amounts of expenditure spent on biomedical research, an especially distinctive feature of the contemporary context is the increasing commercialization of knowledge.

9. Conclusion

In the introduction to part II it was said that the reason for the recent increase in the FDA drug approval rate might be in the strengthened academia-industry relation (Takebe, Imai, and Ono 2018; Robinson 2019) resulting at least partly from the efforts to translate medical knowledge

to new therapies reaching patients (NIH 2014). Translational medicine is dedicated to developing medical products more efficiently, which involves a coordination of research efforts through cooperative work of many stakeholders whose values are often conflicted. In order to evaluate the translational paradigm from an epistemological and ethical perspective, I started my analysis from the assumption that translational approach involves new research methods to achieve the desired goals of acceleration, efficiency, and cost-effectiveness of health care delivery. However, accelerated translations are mostly achieved through an improved exchange between key stakeholders, but also an advancement in basic science, clinical assessment, and technology development that is characterized by a reorientation to personalized therapies. I have pointed out some limitations to the acceleration of biomedical research on the level of drug discovery and on the level of effectiveness assessment, as well as measures that can contribute to acceleration in early drug discovery. I have argued that translational medicine is a structured effort to foster serendipitous discoveries, following an understanding of serendipity as advanced by Copeland (2015, 2019). Also, that contemporary translations – accelerated, efficient, and cost-effective – mostly happen in the domain of personalized medicine.

Taking acceleration as the starting point, I have identified certain epistemic practices that make a difference with regard to speed of research, namely, adaptive design clinical trials on the level of clinical research.⁷³ Through a comparison of randomized controlled trials and adaptive trials I have showed how goals and methods of clinical research can be diverse depending on whether the primary motivation is to understand or to intervene, and whether the goal is to develop products for a larger population of patients or products for individualized care. Randomized controlled trials (RCTs) are the established method in clinical assessment of medical therapies. They present the standard of reliability, but nonetheless face problems of internal and external validity. Moreover, they are often considered too long and inefficient, since the majority of therapies still fail to reach the market. Adaptive trials which can shorten the overall time of clinical assessment are considered less reliable, but justified in cases when quick action is needed and when molecular basis of intervention is well understood. Nonetheless, they also face problems of validity. In personalized medicine, however, a therapy is designed to target a

⁷³ And high-throughput screening and machine-learning as methods used in drug discovery for personalized therapies. However, this was not argued for and should be left for assessment at another place. Though machine-learning methods have been introduced in both case studies and they are related to the topic of time and limited resources, their more detailed analysis and evaluation is outside the scope of this thesis. For my account it suffices to recognize that they contribute to speed of practices.

more constrained group of patients so a very broad generalizability of findings is not required. Furthermore, many personalized therapeutic and diagnostic tools are genetic tests whose risks and benefits are not completely known. Finally, the access to these treatments is currently very restrained because of their cost.

In the case of medical research, the relation between speed and reliability is very complex. The longer the trial takes, the more evident the effects of a therapy are, but collecting more evidence has to be stopped at a certain point in order to make a decision about the efficacy of a new intervention. Often the patients and the producers share an interest in having therapies approved sooner, but there is always a risk involved with regard to appropriate length of a trial. Interventions differ in the speed of their effects and conditions differ in the speed of their symptom development and manifestation. Thus, the risk is either that the trial is too short, hence effects are not yet manifested and the cost is connected to under-evaluation, or that the trial is too long, hence the prolonged trial is unnecessary and the cost is related to over-evaluation.

Inductive risk assessment, however, is very different in the case of medicine than in the case of gravitational wave physics discussed in the first part. In the medical context the acceptance or rejection of a hypothesis about effectiveness of a therapeutic, preventive, or diagnostic product influences patients' wellbeing, but also financial interests of the producer of a therapy. Moreover, results have to be extrapolated from an experimental setting to a real-world context, and many factors may influence the final outcome there. In the gravitational wave physics case the initial preference was for erring on the side of false negatives, i.e. avoiding false positives. Contrary to that, Stegenga (2017) argues that the FDA standard for evaluating medical therapies has a preference for errors on the side of false positives and thus avoids false negatives, namely by applying the 0.05 standard of significance (1 in 20 chances of error) that medications have to satisfy in order to be considered effective. FDA requires two RCTs with a positive result in order to approve a therapy, although it allows for exceptions. Stegenga argues that it is easy to satisfy the 0.05 significance requirement and thus "show" that a medication is effective, even though it is not. The first problem he identifies with FDA regulation is that studies are performed by the producers. Although they have to follow regulatory guidelines (requiring, for example, randomization), there is still a lot of "latitude in how studies are designed, executed, and analyzed" (Stegenga 2017, 23), which allows the intrusion of bias. There is a further problem related to satisfying the p-value standard:

“A more concrete problem with the FDA standard for drug approval is that a standard based on statistical significance lends itself to ‘p-hacking’. Spurious correlations can occur by chance, and the more complex a data set is, and the more analyses performed on a data set, the more likely it is that one will discover a spurious correlation. P-hacking can occur when a researcher exercises ‘researcher degree of freedom’: researchers perform multiple studies, on multiple parameters, choosing which parameters to measure and which comparisons to make and which analyses to perform, and they can do this until they find a low enough p value to satisfy the standard of statistical significance even when the experimental drug is not in fact beneficial.” (Stegenga 2017, 23-24)

Torsten Wilholt (2009) gives further examples of bias towards a preferred outcome (preference bias) entering research design, interpretation of outcomes, and communication and dissemination of results, such as the use of substandard comparison (a placebo instead of state-of-the-art-therapy), or selective reporting of outcomes. Sergio Sismondo has written extensively on the problem of exploiting scientific authorship by pharma industry via the practice of ghost authorship – using academic authors as “celebrity sponsors” (Sismondo 2004, 152) for already in-house written articles which favor new pharma products, in order to acquire credibility for their claims. Another way in which preference bias can influence a positive evaluation of a therapy is experimenting on a patient group which is not adequately representing the target group for the therapy. For example, testing a therapy based on its effect on young people in good physical condition and general health, when the actual, real-world context would be better represented by older people with chronic conditions who will be the likely consumers of the product. Such impermissible practices, coupled with, according to Stegenga (2017), low FDA standard, make it more likely that new therapies are underregulated than overregulated. This means that ineffective and possibly unsafe therapies can easily reach the market, even though most of the therapies still fail to reach the market at all. Taken the interests of industry to recover the costs and their increasing influence on academic institutions as partners in research and development, we can recognize a dangerous leeway for the intrusion of problematic values on research practices in both industrial and academic contexts.

There are opposing views about FDA regulation, however. Julian Reiss (2017) is rather concerned about FDA overregulation which is, according to him, preferred to underregulation because of the risk of harm if unsafe medications reach the market. According to his view, FDA standard is set with a preference to err on the side of false negatives, and the problem is thus in

too strict FDA standards which contribute to possibly effective medications not reaching the patients.

“A beneficial drug that doesn't get approved at all will probably never make the news. A beneficial drug that gets approved later than it could have been will likely be presented as a success story by focusing on the benefits it brings now and not on the harms that were caused by not bringing it to the market earlier.” (Reiss 2017, 167)

“Because of the lengthy testing process the FDA requires, many effective drugs reach the market much later than they could. This also creates harms: patients who die or suffer unnecessarily because essential drugs do not reach them, or do not reach them soon enough.” (ibid., p. 169)

I do not attempt to adjudicate between opposing views of Stegenga and Reiss. The ultimate answer about which one of them is correct requires more empirical research. Still, it is important to see how inductive risk considerations are burdened with much greater risks on both sides of errors in case of effectiveness assessment for medical therapies than in evaluating candidate events for gravitational waves. The standard of significance is also remarkably different: 5 sigma as opposed to 0.05 significance level. Weaker evidence will satisfy when the cost of not doing anything is high. On the other hand, research can last longer if there are no harms done by the absence of its results. Nonetheless, the adequacy of both of these significance levels has been brought into question (Lyons 2013; Stegenga 2017).

The overregulation vs. underregulation controversy has been continuously evoked throughout the discussion of translational medicine case-study because of the inevitable reliance on hindsight knowledge when it comes to favoring speed in effectiveness assessment: we want it when it brings about the better and we do not want it when it brings about the worse, and it is a leap of faith on which side we end up being. Harm is made both if trials are too short and approved medications later turn out to be ineffective and/or unsafe and if trials are too long and approved medications later turn out to be effective and safe. Patients are justified both in wanting access to experimental treatments for life-saving conditions and in wanting better regulation of new therapies. My commitments go only as far as stating the descriptive claim that time is an inevitable dimension of research and that epistemic practices are intertwined with pragmatic considerations about the limitations of resources and how to address them most successfully.

A note is needed on adaptive trials regulation. In October 2018 FDA issued a draft guidance document for industry called *Adaptive Design Clinical Trials for Drugs and Biologics* (FDA 2018e). In this document particular attention has been given to the problem of “Type I error probability inflation” (ibid., p. 8), i.e. inflation of false positives.

“For example, there are a number of ways in which adaptive features can inflate the Type I error probability of a trial. The most obvious examples of this are cases in which multiple statistical hypothesis tests are performed.” (ibid., p. 7)

A further worry is related to reliability:

“Some adaptive design features can lead to statistical bias in the estimation of treatment effects and related quantities.” (ibid., p. 8)

The document gives guidelines on appropriate ways to address and communicate such problems, but I believe these points raise sufficient flags with regard to necessary scrutiny with which adaptive designs have to be evaluated, inside and outside of personalized medicine context. An explicit recognition of the problem of false positive inflation makes a case for Stegenga’s view. However, since false negatives are less often recognized as harms, Reiss’ makes a very good point. It seems to be the case that harms caused by not doing or doing something too late are considered to be less problematic than harms caused by actively doing something. In other words, when a new medication is not effective and especially if it is harmful, this is very directly recognized as ethically problematic, while when a harmful condition is not addressed (soon enough) there is no responsibility attributed to the producer or the regulator for not providing a treatment for the condition on time. Again, it is very hard to discuss the trade-off because it is very hard to weigh and compare harms. A shared premise, however, in both Stegenga’s and Reiss’ account is that effective treatments do not reach the patients, for one reason or another. Stegenga focuses on ineffective treatments, while Reiss focuses on unapproved treatments, and in both cases the patients are at a loss.

Policy response to this widely shared recognition is acceleration – first through translational programs, and recently through personalized medicine. I have argued that the “personalization” of medicine presents the internal change in contemporary biomedical research, and showed how translational medicine has facilitated the transition towards individualized approach. I have also argued that translational efforts have been mostly focused on external measures of addressing the causes of translational gaps, but that it is hard to delineate them from internal measures. In order to adequately assess the role of translational medicine, I have introduced personalized and

evidence-based medicine, which enabled a comparison. The focus has mostly been on the close relation between translational and personalized efforts which share a common goal to accelerate the delivery of health care and to empower community and patients as partners in research, but also to address the “pipeline problem” of medical innovation.

Evidence-based medicine (EBM), on the other hand, has been important for this analysis because of the emphasis on RCTs and meta-studies as sources of high evidence. I have mostly followed Miriam Solomon’s (2011, 2015) treatment of EBM since she explicitly brings the emergence of translational medicine in connection with certain shortcomings of the EBM paradigm. It is worth noting Mark Robinson’s (2019) follow up on Maienschein’s et al. (2008) and especially Solomon’s (2015) treatment of translational medicine. He reflects on why it is impossible to find solid and coherent epistemological input in translational medicine (notably, fourteen years after its seminal programmatic document). Robinson identifies a common problem in all attempts to philosophically analyze new medical models, particularly translational medicine, and that is a striking lack of novelty and epistemic substance. He asserts that traditional epistemology is not equipped with sufficient tools to analyze the phenomena of new medical movements, and that insights from political economy of research and innovation might shed better light on the broader socio-economic context in which translational medicine emerged. Moreover, that the apparent lack of epistemological input does not mean that translational shift is not influencing knowledge production. He therefore calls for more social epistemological work which is “uniquely positioned to consider the impacts of changing ‘epistemic systems’” (Robinson 2019, 3). Robinson’s summary of the primary function of translational medicine hits the target very precisely:

“In other words, one of the primary functions of TrM is as a means through which private firms can externalize the risks and costs of early stage biopharmaceutical R&D. (...) TrM must be analyzed in relation to its financial functionality – a key cog in a larger global strategy regarding R&D innovation. If we take this view, it will therefore make sense that recent analyses of TrM reveal little that is particularly epistemically novel. (...) Through TrM, newly risk-anxious industry partners could both manage portfolio risk, and retain rights to monetize university partners’ potential scientific findings. Thus, while TrM funding brought about a rapid expansion of centers, doctoral programs, faculty lines and scientific conferences, it also brought with it a key infrastructure to insert discrete industry interests into university science projects. It turned university TrM

centers into biotech startups and pharmaceutical partners into a class of investors.”
(ibid.)

I hope that this relation has become clear in my analysis and that my treatment of translational medicine has been inclusive enough to shed light on the complex interactions at work. My attempt was to evaluate translational medicine as a declarative and actual scientific resource of epistemic and non-epistemic tools that accelerate the achievement of epistemic and non-epistemic aims. It was necessary to involve many insights that transgress the boundaries of disciplines and evidently involve not only a social epistemic, but also an ethical dimension. In translational medicine, considerations that deal with temporal aspects of research (When to stop a trial? How to accelerate research?) are considerations about harms, either harms caused by not translating knowledge to practice, or harms caused by translating it badly.

This analysis has been done as a part of a broader aim to confront two practices from the opposing ends of the so called basic-applied research continuum. While even the gravitational wave case offered a less than a crystal-clear epistemic picture in which the distinction between inter-scientific and extra-scientific factors got blurred, the translational medicine case has definitely erased the boundaries that might have still been lurking in the back. The transition, however, has been gradual. I have in both cases started with internal-external distinctions for the sake of clarity and continuity with the scholarly tradition, but have ultimately abandoned such conceptualization. The disappearance of boundaries is consistent with the thesis I have set out to advance: that epistemic aims are pragmatically informed, and that epistemic values involve a pragmatic dimension, since they promote the attainment of truth in constrained circumstances of limited resources. Through the discussion of case studies we have seen how different epistemic aims prioritize different methods. We have seen how speed is highly valued, yet brings epistemic and non-epistemic risks which need to be accounted for. We have also seen how epistemic and non-epistemic value conflicts (mutual and internal) permeate the research process. In the next, final part of the thesis the focus will be on identifying the relation between time, science, and values in more detail. Concepts like “epistemic”, “non-epistemic”, and “pragmatic” have been used several times in the analysis of both case studies, leaving many questions open with regard to their precise meanings. Part III will provide a clarification.

Part III

Time, Science, and Values

1. Introduction

The topic of values in science is one of the most fruitful and rapidly evolving areas of philosophy of science. It is focused on explaining the choice of theories and methods, hypothesis acceptance, and model building, in terms of characteristics that these representations possess in order to be preferred in belief and action. Epistemic aim in terms of belief is to form true conclusions, hence the epistemic values are those indicative of truth. Epistemic aims in terms of action can be diverse: to model (a waveform), to predict (the effectiveness of an intervention), to assess (a risk), to detect/observe (a gravitational wave). Epistemic values are also indispensable for the fulfillment of non-epistemic aims: to intervene (on the cause or the symptoms of a disease or a condition), to regulate (a new substance in use), or to decide policies on various issues (climate, health, education, or economy).

Case studies discussed in this thesis are chosen with the aim to represent scientific practices from the opposing ends of the basic-applied research continuum, although the distinction is a controversial one and I have not committed myself to resolving its problems. Normally, the aim of basic science is to acquire understanding and knowledge, i.e. true beliefs. Epistemic values are thus those indicative of a truthful representation or explanation of the observable phenomena. Gravitational wave physics that is supposed to be characteristic of this kind of research involves an observation (or a detection) of a predicted but previously yet unobserved phenomenon. What has to be established with certainty is a true belief about the observation. The underlying theory, Einstein's theory of general relativity, had already been decisively confirmed based on observational and experimental evidence, and the project's goal had been to directly observe gravitational waves which the theory predicts. What had to be ascertained is the confidence that the observation was achieved, and it was important to avoid a false belief about the first gravitational wave detection event. Two reasons have been given for this. The first is the epistemic reason – building knowledge on a false fact would lead to further mistakes. The other reason appears to be non-epistemic – a false announcement of a detection would undermine the trust in project. But this would again lead to negative epistemic consequences –

without governmental support, the investments necessary to pursue the project might be decreased, and thus less knowledge would be acquired.

The other case study, on translational medicine, is to a lesser degree related to the epistemic aim of understanding and acquiring knowledge about the mechanisms of intervention. It is rather related to the aim of assessing effectiveness in order to, ultimately, intervene on causes or symptoms of a disease or a condition, but also to inform the regulators and to bring a therapeutic product on the market. It has been shown that assessing effectiveness is often influenced by market interests and urgent need to address problems of societal importance. In the trade-off between stopping a clinical trial early, possibly before the onset of certain effects, and stopping a trial late and thereby risking the unnecessary cost of prolonged experimentation, values which are not primarily related to the pursuit of truth, but rather at lowering costs or providing earlier access to treatments, influence the decision. Unfortunately, what follows is often a subsequent recognition of false beliefs about efficacy and safety, i.e. a false positive. However, a false negative is also a mistake which involves harm. In this case it is harm caused by not addressing a health condition by failing to translate knowledge into medical therapy or failing to translate it quickly enough.

The other reason for the choice of case studies have been their differing commitments with regard to time it takes to fulfill research aims. LIGO research is a long-term effort focused clearly on one epistemic goal, which is nowadays rapidly changing in the face of new detections. Translational medicine, on the other hand, is explicitly dedicated to accelerating discovery and research, with programmatic documents stating 5-10 years goals. The preference for such a short temporal range has been critically assessed, especially because of evidence that shorter research and evaluation time often compromises the safety of patients. However, a recent turn to personalized therapies has brought about a change in the discovery and clinical assessment route, in which adaptive studies have started to replace the longer, though arguably more reliable randomized controlled trials. Better clinical outcomes of personalized therapies are expected because of a holistic approach to patients which is hoped to provide more predictive effects.

A trade-off is evident: between understanding (RCTs) and intervening (adaptive trials) as exemplified in the Ebola and AZT case which are characterized by high social stakes and urgent need for action, and between generalizable findings (RCTs) of EBM and individualized therapies (adaptive trials) characteristic of personalized medicine approach. The aims are traded

off against one another through a different prioritization of epistemic values that the competing methods are comparatively characterized by. In the cases of social urgency (Ebola and AZT), reliability of RCTs is traded off against the speed of adaptive trials (Ebola ca sufit!). In the second case (personalized medicine), the scope and robustness of RCTs is traded off against precision and flexibility of the personalized approach as operationalized through adaptive design trials. Furthermore, trade-offs that hinge on considerations about time through attempts to increase the speed of methods are also characteristic of gravitational wave research, which is not burdened by direct social stakes. It has been shown how considerations about time and speed nonetheless permeate every stage of the research process. Most explicitly, methods in waveform modeling involve a trade-off between speed and reliability.

Apart from the particularities of these cases, more general arguments can be made about the relation between epistemic, non-epistemic, and pragmatic values, three categories used throughout the analysis of case studies. In order to take up this task more directly, a brief overview of these concepts and accounts that employ them is in place. This will be given in Chapters 2 and 3. My ultimate answer about the relation between these different categories of values will need to be consistent with the already established erosion of the distinction between what is internal and what is external to science in both of the case studies discussed. I will subscribe to the view of Matthew Brown (2013) and reject the “lexical priority of evidence over values”, in order to bring the epistemic, non-epistemic, and pragmatic distinction on more even grounds, namely, by pointing to the pragmatic dimension that epistemic activities have in their pursuit of true conclusions. Since these pragmatic aspects are often supported by non-epistemic considerations, a commitment to a sharp distinction will be very hard to uphold. However, I will not attempt to take part in the debate about the legitimate role for non-epistemic values in theory assessment, because this would oblige me to go much deeper in the debate between the proponents of the so called “aims approach” (Elliott and McKaughan 2014; Intemann 2015) and the “epistemic priority approach” (Steel 2010, 2016) than I should go if I want to keep the focus on aspects of time as a limit of resources. I will only criticize certain aspects of Steel’s “epistemic priority approach”, which is consistent with subscribing to Brown’s rejection of lexical priority of evidence over values. My overall stance will support a broader, inclusive, and functional role for non-epistemic values in the internal stages of scientific research. This support clearly stems from the fact that the internal-external distinction has gradually disappeared in both of my case studies.

I will start by giving an overview of standard accounts of epistemic, non-epistemic, and pragmatic values, and their often conflicting roles in decision making in science. I will introduce main arguments that employ these concepts, such as the inductive risk argument which requires a distinction between epistemic and non-epistemic values, and the underdetermination argument which does not require such distinction.⁷⁴ I will then focus particularly on the role of time-sensitivity in science by drawing on a discussion between Kevin Elliott and Daniel McKaughan (2014), and Daniel Steel (2016). They discuss the role of non-epistemic values in theory assessment and the epistemic status of speed of practices to generate results. I will argue that speed takes priority over ease of use in the cases they discuss and that speed should be considered an epistemic value, but also that Steel's account of epistemic values (2010) bears a commitment to particularism about what is epistemic and non-epistemic in individual instances of research. Finally, I will give an account of time-sensitivity in science and relate it to cases of transient underdetermination with the help of case studies discussed in the first two parts of the thesis. I will take a stance that the fulfilment of epistemic aims is inherently pragmatic and that non-epistemic values also have a functional role which should not be downplayed, though I will not explicitly argue for the latter claim. My overall account will nonetheless maintain that the influence of all three kinds of values is necessary for socially responsible scientific work. Moreover, that the borders between these categories of values are often hard to draw.

⁷⁴ Although Heather Douglas (2009), who is one of the most prominent defenders of the argument from inductive risk, replaces the epistemic/non-epistemic value distinction with epistemic criteria instead of epistemic values, and introduces indirect and direct roles for values in order to specify when their influence is legitimate and when it is not, her account still needs a clear distinction between what is epistemically acceptable and epistemically unacceptable (see Steel 2010). De Melo-Martín and Intemann (2016) even argue that the inductive risk argument does not really challenge the value-free ideal since it does not allow that non-epistemic values play a positive role in determining evidence. Douglas herself admits that “a clear demarcation between epistemic (acceptable) and non-epistemic (unacceptable) values is crucial for the value-free ideal” (2009, 89-90). Since her account retains a version of this distinction, her challenge to the value-free ideal is questioned by the proponents of the descriptive claim about the role of values in science. This claim states that non-epistemic values necessarily fill in the gap between theory and evidence since evidence is always inconclusive. Proponents of the descriptive claim also hold that non-epistemic values can play a positive role in determining evidence, and that the distinction between epistemic and non-epistemic values should therefore be questioned (Longino 1996; Rooney 2017).

2. Epistemic and non-epistemic values

The role of values in science, especially the role of non-epistemic values, is a longstanding topic of discussion in philosophy of science. Francis Bacon's idols are a good historical point of departure when engaging in it. In his *New Organon* published in 1620 he wrote about four kinds of idols that obstruct human reasoning: idols of the tribe (human limitations, perception), of the cave (custom, upbringing), of the market (communication), and of the theater (ideologies, philosophical theories) (Bacon 2009 [1620]). Social, political, moral, and ideological values present a contemporary counterpart to Bacon's idols, often depicted under the concept of non-epistemic values. They are characterized by the perceived flaw of obstructing the attainment of truth, hence posing a concern for the value-free ideal, an ideal rooted in the view of science comprised of facts only. This ideal has in the last few decades been widely recognized as unattainable due to many evaluative judgments that permeate scientific research, especially when deciding on which lines of research to pursue and how to apply the acquired knowledge, but most often even unattainable in the process of scientific justification itself.

The value-free ideal has been most substantially brought into question by the inductive risk argument, the underdetermination thesis, and the recognition of the unreliability of the epistemic/non-epistemic value dichotomy. A standard version of the inductive risk argument was formulated in 1953 by Richard Rudner, stating that scientists have to take into account the outcomes of accepting false hypotheses, especially when the consequences of such acceptance bear moral weight, i.e. lead to detrimental health effects or other significant losses which are morally inadmissible.⁷⁵ If false positives lead to negative non-epistemic consequences, researchers are justified in adjusting evidential standards in a way that avoids false positives and tolerates false negatives, rather than the other way around. This means that the standards for acceptance of hypotheses in these cases have to be higher. Such considerations mostly occur in applied contexts, while basic science was still somewhat intact by the inductive risk argument, since its primary aim is knowledge acquisition and not action with social consequences, such as regulation of chemicals or therapeutic intervention.

Apart from Rudner's (1953) formulation, an influential contribution to the inductive risk discussion is Hempel's (1965), while two standard responses to Rudner are those by Jeffrey (1956) and Levi (1960). Jeffrey (1956) criticized the assumption that scientists should make

⁷⁵ The problem of inductive risk was introduced in section 5.2 of part I, p. 53-55.

value judgments and contended that their role is to assign probabilities to hypotheses, while the policy makers' job is to make value judgments on the basis of assigned probabilities. Steele (2012) gives a compelling response to this argument:

“In practice, scientists will simply make the judgment themselves about whether more evidence would be useful. In effect, this is another stronger, pragmatic argument for why scientists must make value judgments. It is one thing for policy makers to be informed of the subtleties of the scientist's beliefs regarding states; it is another thing for policy makers to also be informed of how the scientist's beliefs may change given the various types of evidence that may be collected in the future. (...) Once we consider the scientist's role in initially specifying a policy decision, as well as when to actually make the decision versus collect more evidence, we see how entrenched the scientist is in social-political value territory.” (Steele 2012, 903)

Pragmatic aspects of scientific research such as the one Steele puts forward, stemming from temporal limitations and the need to make decisions, are also in the focus of this thesis and its two case studies. Decisions influenced by temporal limitations are not only made on the level of hypothesis acceptance, but also on the level of goal setting, choice of methods, and communication of results.

Levi (1960), on the other hand, responds to Rudner (1953) by arguing in favor of resolving the uncertainty by epistemic values only, namely by increasing the evidential power. However, Staley's (2017) interpretation offers a less value-free reading of Levi's account:

“On Levi's account, value judgments of many kinds might (in their 'indirect role?') contribute to a decision to let one's inferences be governed by a determinate degree of caution. But the autonomy of the epistemic is preserved insofar as the inferences thus carried out can be understood strictly as attempts to replace agnosticism with belief, such that correct answers are preferred to incorrect answers, and such that the decision whether to accept a hypothesis as an answer rather than suspend judgment is based on a balance between an interest in relieving agnosticism and a cautiousness regarding the risk of embracing the answer erroneously.” (Staley 2017, 48)

We could see how the gravitational wave researchers were balancing their interest in relieving agnosticism against cautiousness regarding the risk of falsely accepting a hypothesis about the first detection by setting and, ultimately, satisfying a very high statistical significance level. It was because their interest is not only true belief, but also a high degree of confidence that it is

a true belief, so that the detection can be efficiently communicated to the public. Initial preference for erring on the side of false negatives, and the need to evaluate the candidate event for months until a high significance level has been confirmed, were not only present as a consequence of epistemic interest but also as a consequence of non-epistemic interest. As I emphasized in the discussion of Part I, this makes the inductive risk argument valid also in basic contexts, and not only when applied to considerations of consequences in applied research. Evidential standards are chosen based on non-epistemic considerations both in gravitational wave research and in biomedical research, partly because both cases involve an action in addition to belief: to communicate the detection in the former case, and to intervene in the latter. Rudner's original argument has been widely discussed and eventually broadened, especially by Heather Douglas (2000, 2009), to establish a normative stance about the role of values in theory assessment: they should influence scientific research when hypothesis acceptance or rejection leads to morally relevant consequences in different stages of scientific work: choice of methodology, characterization of data, and interpretation of results. Douglas proposes a distinction between a direct and an indirect role for values,⁷⁶ in order to give an account of when the influence of non-epistemic values is permissible and when it is not. It is permissible only as an influence on the choice of evidential standards for accepting a hypothesis, in other words, non-epistemic values are allowed to have only an indirect role. When values are direct reasons for accepting a hypothesis, such influence, according to Douglas, is impermissible. It is justified to have a higher standard for evaluating hypotheses whose acceptance might have ethical consequences, than for evaluating those that do not have such consequences. In a purely epistemic sense, there should be no difference in evidential standards among different cases, if the only goal is truth. But if the goal is also to prevent harm, then potentially harmful consequences require stricter evidential standards. This relation also involves a temporal dimension.

“In some cases (e.g., detecting the Higgs boson), waiting until you have really strong evidence does not entail any downsides – there are no practical uses (as yet) for the particle, and so waiting for strong evidence before confirming its existence (and avoiding strongly the false positive) is fine. In other cases (e.g., detecting an emerging disease with pandemic potential), weaker evidence will suffice before sending up

⁷⁶ Which Staley refers to in the previous quote.

warning flags, as there are real downsides to waiting too long before making public claims (the false negative risks are very potent).” (Douglas 2017, 89)

Such an underdetermined relation between hypotheses and evidence is an inevitable feature of scientific practice and it has been in our focus in the discussion of two case studies. It is the relation of transient underdetermination of theories by evidence, or temporary underdetermination by currently available evidence (Kitcher 2001; Biddle 2013b). Since waiting indefinitely for sufficient evidence is rarely ever an option, decisions have to be made in the face of uncertainty and non-epistemic values have to fill in the gap between evidence and hypothesis assessment (Biddle 2013b). This is especially so when non-epistemic stakes are high, possibly urgent, which means that not acting would lead to a loss of life, health, or other value, such as biodiversity, social stability, or wealth. Moreover, we have seen that waiting indefinitely is not an option also in basic science such as gravitational wave physics.

The impression of an apparent lack of rationality in theory change that emerged after Thomas Kuhn’s *The Structure of Scientific Revolutions* in 1962 boosted the belief that social, moral, and political values play a role in theory acceptance, or at least have done so in the past. In response to this unwanted impression, Kuhn established an influential concept of epistemic values in his 1977 paper “Objectivity, Value Judgment, and Theory Choice”, arguing that these values particularly are the ones that guide decision making in science. Kuhn used the term “epistemic values” for five characteristics of theories: accuracy, consistency, scope, simplicity, and fruitfulness (1989, 357). *Accuracy* means that the consequences of a theory are “in demonstrated agreement with the results of existing experiments and observations” (ibid.). *Consistency* can be internal and external: propositions of a theory should include no contradictions within themselves (internal) and with other accepted theories (external). Broad *scope* of a theory presupposes that its consequences “extend far beyond the particular observations, laws, or subtheories it was initially designed to explain” (ibid.). *Simplicity* of a theory is for Kuhn understood as “bringing order to phenomena that in its absence would be individually isolated and, as a set, confused” (ibid.). *Fruitfulness* means that a theory should “disclose new phenomena or previously unnoted relationships” (ibid.). These “criteria of choice” function “not as rules, which determine choice, but as values, which influence it” (ibid., p. 362). Kuhn argues that they cannot be applied algorithmically and that there is, first, always a subjective judgment involved in identifying whether a certain value is instantiated in a theory, and second, always a subjective preference involved in ascribing relative weight to one value at the expense of other.

Informed by the case studies, we can take the example of scope vs. accuracy (in our cases applied to methods, not to theories). In order to establish a parameter space of a broad range of waveforms, gravitational wave researchers have to use a faster but less accurate method. If they want to model fewer waveforms in more detail, they have to sacrifice the aim for establishing a parameter space. In the medicine case, the efficacy assessed in RCTs often fails to extrapolate to effectiveness in real world settings because of many diverse factors influencing the outcome, while in individualized approach a more predictive outcome is expected, but at the expense of establishing effectiveness more broadly than for a constrained set of patients.

Establishing a category of epistemic values made the longstanding value-free ideal look quite wrong, when there obviously are values that science relies on, namely, epistemic values. The new reading of the value-free ideal now banned only the influence of non-epistemic values on the research process. It did not take long to notice that non-epistemic values affect scientific research whether we want them there or not, as it became obvious from historical evidence that science has often been infused with political views and ideology, for example about the subordination of women or the superiority of the white race. Feminist epistemologists have from the 1980s on taken up the older Duhem-Quine thesis about the underdetermination of theory by evidence and brought a new awareness of the evidence-theory gap which is often filled by non-epistemic values and background assumptions (Longino 1990, 2002, 2004). For example, by internalized ideological presuppositions that do not have to be made explicit. This is the central premise of the underdetermination argument: since every theory is underdetermined by available evidence and logic, background assumptions and non-epistemic values fill the gap between evidence and theory. It is also a descriptive claim about the role of values in theory and hypothesis assessment: regardless of whether they should be there or not, we necessarily find them there. This is made explicit by the already mentioned Biddle's (2013b) argument from transient underdetermination, of which more will be said in Section 4.5.

Kuhn's list of epistemic values has been reassessed many times, and different values have been added or subtracted from it, for example, simplicity and explanatory power. A case has been made by Longino (1995) that simplicity can only be an epistemic value if the world is indeed simple, and that many false theories can satisfy this requirement. A similar argument goes for explanatory power: many theories that turned out to be false were nonetheless able to explain many phenomena. Besides suggested subtractions from the list, new epistemic values have been proposed. For example, Longino (1995) argues for diversity as an epistemic value, since a plurality of perspectives broadens knowledge, sheds light on biases, and thus leads to a more

objective account of phenomena. Epistemic values that usually survive on all the lists are the ones indicative of truth in all circumstances: empirical adequacy and internal consistency. There are also arguments against upholding the distinction between epistemic and non-epistemic values, stemming primarily from the underdetermination argument. Since there is no agreement on what should count as an epistemic value and we know that different kinds of values permeate the research process, the distinction should be abandoned (Longino 1996; Rooney 2017).

According to Longino's contextual empiricist account (1990, 2002), objectivity can only be attained by mutual criticism, i.e. the quest for objectivity is only possible as a social goal, not as an individual undertaking. Many background assumptions, individual and group biases (in the vein of Bacon's idols), and idiosyncrasies are impossible to be individually identified and balanced out. These values have to be made explicit and communicated in order to be evaluated and taken up by others in a community based search for knowledge.

“The critical dimension of cognition is a social dimension, requiring the participation of multiple points of view to insure that the hypotheses accepted by a community do not represent someone's idiosyncratic interpretation of observational or experimental data.”

(Longino 2002, 106)

Alvin Goldman's social veritistic epistemology (1999) also takes the acquisition of knowledge to be a social endeavor, in addition to individual belief acquisition. However, accounts such as Longino's (1990, 2002, 2004), Rooney's (2017), Rolin's (2017), and de Melo-Martín and Intemann's (2018) put forward the view that science is not only necessarily social, but also that the influence of values on research is positive, including the influence of non-epistemic values. The ideal that Longino puts forward instead of the value-free ideal is the social value management ideal, according to which values are collectively managed in social interaction, but only when scientific community satisfies four criteria: there have to exist publicly recognized venues for criticism, participants have to be responsive to criticism, they have to share the same standards, and tempered equality of intellectual authority has to be granted (Longino 1990, 76-81; 2002, 129-131). Such value-laden account of science is summarized by Longino in three related points:

“1. Value-ladenness does not mean that social values outweigh other considerations, but that they interact with data and hypotheses in determining evidential relevance.

2. If the value-ladenness of a theory is a contingent matter, then whether values are extricable or inextricable from inquiry is not a matter for general a priori argument but case-by case-analysis.

3. Contextual or social values are not just negative features in inquiry, but can have a positive role in grounding criticism of background assumptions and in fostering the development of empirical investigation in directions it would not otherwise go.” (Longino 2002, 51)

3. Pragmatic values

Pragmatic values are the least recognized and used category of values, though they have been acknowledged in at least two accounts of values in science, those of Ernan McMullin (1982) and Heather Douglas (2013). In addition to their accounts, I will briefly present Matthew Brown’s (2013) account of pragmatist functionalism about inquiry which supports the view of pragmatically informed pursuit of epistemic aims – a knowledge and action-driven pursuit informed about the limitations of time, material, cognitive and computing resources, as exemplified in the two case studies in this thesis. Brown’s account does not establish a pragmatic category of values, but it recognizes that the pursuit and fulfillment of epistemic aims involves a pragmatic dimension. This is made explicit in Brown’s argument against the lexical priority of evidence over values, which allows a functional role for evidence and values equally. I argue that a part of this shared functional role is to address the limitations of resources, which makes Brown’s account supportive of my view that pragmatic factors are, ultimately, epistemic.

3.1 Ernan McMullin’s account

In Ernan McMullin’s Presidential Address “Values in Science” at the 1982 Philosophy of Science Association meeting, pragmatic values are added alongside epistemic and non-epistemic values.

“When I say that science is value-laden, I would not want it to be thought that these values derive from theory-appraisal only. Value-judgment permeates the work of science as a whole, from the decision to allow a particular experimental result to count as “basic” or “accepted” (the decisional element that Popper stressed), to the decision

not to seek an alternative to a theory which so far has proved satisfactory. Such values as these may be pragmatic rather than epistemic; they may derive from the finiteness of the time or resources available to the experimenter, for example. And sometimes the borderline between the epistemic and pragmatic may be hard to draw since (as Duhem and Popper among others have made clear) it is essential to the process of science that pragmatic decisions be made, on the temporary suspension of further testing, for example.” (McMullin 1982, 18)

Exactly these kinds of decisions, related to the limits of time as a resource, have to be made in both of our case studies. Gravitational wave researchers have to be fast in evaluating candidate detection events and biomedical researchers have to stop the trials in order to reach a conclusion, thereby risking possibly high social stakes either from prolonged experimenting or from prematurely accepting a hypothesis. A paradigmatic example of McMullin’s pragmatic value, stemming from limited temporal resources, would thus be speed.⁷⁷ Speed is a value instantiated in practices, methods, and processes, not theories and hypotheses. A theory cannot be faster than another theory, but it can possess certain values that contribute to speed of generating results, such as ease of use, of which more will be said in Chapter 4. However, a method can be fast (post-Newtonian approximation when compared to numerical relativity) and a practice or an experimental process can be fast (adaptive trials when compared to RCTs). Note that the attribution of a certain degree of speed at which methods operate is always possible, but that the evaluation that a method or a practice is fast can only be done in relation to another method which is slower. Furthermore, speed of practices to generate results has to fit a socially relevant time-range, i.e. a method has to operate and produce results in a time frame which is reasonable for researchers to use when working on a particular problem. More precisely, it has to deliver results in a course of years or decades the most, hardly ever a century.

A notable inversion needs to be brought to attention: while the initial setting of the thesis was the one in which gravitational wave research takes “long” and is “slow”, and medical research attempts are taken to be “fast” (Ebola ca suffit!) and shorten the time of research (translational medicine), we could see that the evaluation of the candidate gravitational event lasted five

⁷⁷ I will argue in Chapter 4 that speed should be considered an epistemic value, which is consistent with my view that pragmatic is epistemic since epistemic aims are pragmatically informed. It is also consistent with McMullin’s view that the borderline between the epistemic and pragmatic is hard to draw.

months in the first detection case (exactly from September 14, 2015 to February 11, 2016), while trials in medicine can last for decades. The differences between the two cases are huge, in terms of possible non-epistemic consequences, ethical aspects of research in the case of medicine, different evaluation methods, and most importantly, extrapolation and a degree of generalization that is an objective in clinical trials. The initial contrast between the first detection of gravitational waves and Ebola ca suffit! adaptive trial was thus called “an unfair comparison” not without a reason. The success of the trial is especially spectacular taken that it lasted only few months.

McMullin was also the first (to my knowledge) to use the term non-epistemic values, which he defines in the following way:

“When no sufficient case can be made for saying that the imposition of a particular value on the process of theory choice is likely to improve the epistemic status of the theory, that is, the conformity between theory and world, this value is held to be non-epistemic in the context in question.” (McMullin 1982, 19)

Hence, efficient communication, economic interests, and protection of health are all concerns not directly related to the goal of truth. The fact that we want effective therapies will not make them more effective, neither will the fact that they are on the market. The intrusion of non-epistemic values on the research process is sometimes called “wishful thinking” (Brown 2013).

Note that speed also falls under McMullin’s definition of non-epistemic values since it does not improve the conformity between theory and world. However, speed is not instantiated in theories but in methods and practices, as already said. Methods which are faster may improve the conformity between theory and world indirectly, by providing more time for using and additionally evaluating results which would otherwise have to be waited for. Speed definitely requires a sufficient degree of accuracy, but once this degree is present, faster practices advance the epistemic status of a person or a community that uses them by providing access to results faster, thus enabling more knowledge acquisition on the basis of these results. Moreover, even a very robust epistemic value like internal consistency still requires that the consistent propositions are, ultimately, accurate. The fact that speed does not itself point to accuracy does not mean that it is not epistemically beneficial in the context of inevitable temporal embeddedness of scientific practice. On the contrary.

3.2 Heather Douglas' account

Another account of pragmatic values comes from Heather Douglas (2013). She argues that there are four distinct groups of values that “normally get lumped together” (2013, 796): “values that are minimal criteria for adequate science, applied to the theory per se”; “values that are minimal criteria for adequate science, applied to the theory in relation to evidence” (ibid., p. 799); “values that are desiderata when applied to theories per se”; and “values that are desiderata when applied to theories in relation to evidence” (ibid., p. 800). In the first two groups are values that are “genuinely truth assuring, in the minimal sense that their absence indicates a clear epistemic problem” (ibid., p. 799). For example, internal consistency (of a theory per se) and empirical adequacy (in relation to evidence).

The group of our primary interest is the third group, values that are desiderata when applied to theories per se. These values, Douglas argues, can be considered “strategic or pragmatic values”, but in another place she calls them “cognitive values” (2009, 93-94), identifying them with Steel's (2010) extrinsically epistemic values (Douglas 2013, 800). According to her account, they “aid thinking” so that theories that instantiate them are “easier to work with” (ibid.). For Douglas, these are scope, simplicity and explanatory power.

“These are values that, when instantiated solely by the theory or claim of interest, give no assurance as to whether the claims that instantiate them are true but give us assurance that we are more likely to hone in on the truth with the presence of these values than in their absence. (...) Simpler claims are easier to follow through to their implications. Broadly scoped claims have more arenas (and more diverse areas) of application to see whether they hold. Theories with potential explanatory power have a broad range of possible evidential relations. (...) It is easier to find flaws in the claims and theories that instantiate these values. It is easier to gather potentially challenging (and thus potentially strongly supporting) evidence for them. In this sense, all of these values fall under the rubric of the fruitfulness of a theory.” (ibid.)

Douglas' understanding of pragmatic values is very different from McMullin's and from an understanding that has been in the focus of this dissertation. McMullin's and my account of pragmatic stem from limitations of resources and apply to practices (which can instantiate the value of speed) and decisions (to stop further testing). Douglas' epistemic values apply to theories and are broadly characterized by contributing to ease to use. This understanding of pragmatic is thus not of our interest, since the case studies discussed are primarily about

scientific practices, and not scientific theories. Furthermore, in Chapter 4 I will turn to Steel's account which is much broader than Douglas' account and allows values to be manifested not only by "theories, explanations and hypotheses" (Douglas 2013, 800), but also by "methods, social practices, and community structures" (Steel 2010, 19). Douglas, however, identifies her account of pragmatic values with Steel's extrinsically epistemic category of values, which is a relation that does not seem to hold.

3.3 Matthew Brown's account

The third account focused on pragmatic aspects of decision making in science has been put forward by Matthew Brown (2013) and argues in favor of a pragmatist functionalism about inquiry. His account draws on John Dewey's pragmatism and "differentiates the functional roles of evidence, theory, and values in inquiry" (2013, 837).

"This retains the idea that all three have to be coordinated and that each is revisable in the face of new experience, while introducing further structure into their interactions. According to such an account, not only must evidence, theory, and values fit together in their functional roles, but they must do so in a way that actually resolves the problem that spurred the inquiry." (ibid.)

Brown argues against the lexical priority of evidence over values which he identifies as a shared premise in both of the arguments that support the value-ladenness of science: the underdetermination argument and the inductive risk argument.

"Both arguments begin from a situation where the evidence is fixed and take values to play a role in the space that is left over. The reason that values must play a role is that uncertainty remains once the evidence is in. In a relatively weak version of this argument, social values fill in the space between evidence and theory because something has to, so it might as well be (and often is) social values. In more sophisticated versions, we must use social values to fill the gap because of our general moral obligation to consider the foreseeable consequences of our actions, including the action of accepting a hypothesis. The arguments of these two general forms all assume the lexical priority of evidence over values." (ibid., p. 834)

He further identifies two problems with the lexical priority premise: first, an uncritical stance towards the status of evidence, and second, a reduction of the value judgment idea to a mere

expression of preferences. The reason for the first problem is that “evidence can turn out to be bad in all sorts of ways: unreliable, unrepresentative, noisy, laden with unsuitable concepts and interpretations, or irrelevant for the question at hand” (ibid., p. 836). The reason for the second problem is that we can in fact have good reasons for our value judgments.

“Valuing may be the mere expression of a preference, but value judgments are reflective decisions about values, and properly speaking must be made on the basis of reasons (and judgments can be better or worse because they are made on the basis of good and bad reasons). Value judgments may even be open to a certain sort of empirical test, because they hypothesize relationships between a state or course of action to prefer and pursue and the desirability or value of the consequences of pursuing and attaining them (...).” (ibid.)

Though Brown’s account does not make explicit the category of pragmatic values as such, it does bring a pragmatic stance into focus by allowing both values and evidence to jointly contribute to choices in science. Moreover, his account emphasizes efficiency since the functional roles of values and evidence interact in order to “actually resolve the problem that spurred the inquiry” (Brown 2013, 837). Formulating the epistemic aim in terms of problem-solving acknowledges the pragmatic dimension, as opposed to formulating the epistemic aim in terms of abstract truth-seeking. Limitations of resources are one of the reasons to take such a pragmatic stance.

For example, in gravitational wave research, evidence is noisy (thus also a limited resource) and many uncertainties permeate the research. Speed of waveform modeling is important since it enables researchers to cover a range of waveforms which are indispensable for filtering the data. For now, only the modeled waveforms can recognize a gravitational wave signal in the data, so the more waveforms modeled, the more possibilities for a detection acquired. A parameter space of waveform models has to be established in order to decrease the number of false negatives, and the speed of waveform modeling practices, a value instantiated particularly in the post-Newtonian analytical approximation method, contributes to that. On the other hand, the accuracy of numerical relativity method, which is much slower than the post-Newtonian method, establishes the degree of tolerated error and therefore safeguards the empirical adequacy of the process. By combining the two methods, the modeling process is efficient: sufficiently empirically adequate and timely. Speed is here important both for epistemic and non-epistemic reasons. Empirical adequacy has to be established in a time frame that is socially

acceptable, which means that it has to be achievable in the course of decades rather than in the course of centuries.

In the medical context, evidence of effectiveness is also limited, and there are two ways of making it more predictive: either by narrowing the scope of targeted patients and increasing the evidential power by a variety of evidence, as instantiated in the personalized medicine approach, or by conducting several and longer trials, as instantiated in the EBM approach. Speed is here especially valuable because of non-epistemic reasons, and as Brown rightfully argues, these are very good reasons. Both the need to improve health and to enable cost-effectiveness are valid non-epistemic reasons to think about the time of implementation of medical therapies. Possible conflicts between these two aims are not appropriately resolved by a change in attention to speed, but rather by a change in attention to evidence assessment.

For example, companies may favor shorter trials because of their lower cost, which may be a preference which is in conflict with the safety of patients. The problem in this case is not the value of speed of practices per se, but the inappropriate means of achieving it. The appropriate way to insist on a shorter trial would be to increase the power of evidence so that the outcomes are more predictive, which is a line that personalized medicine has taken. Speed is thus not a problem in itself, the problem is rather in ways of achieving it. The appropriate length of medical trials will always be uncertain, and the trade-off is inevitable: between harms done by early stopping and harms done by prolonged experimenting, which means that increasing speed does not necessarily increase harms. Since the primary goal of biomedical research is to decrease harms by providing new therapies, marrying this aim solely to longer trials is not justified. We want better evidence and predictive outcomes, and this is what personalized medicine is trying to achieve in addition to speed. Contrary to that, stopping a trial early without appropriate evidence of its effectiveness is unjustified, but not solely because of inappropriate length of the trial, but because of inadequate evidence acquired during its course.

The risk will always maintain, regardless of when the trial is stopped and regardless of whether the approach is personalized or not, and sometimes there will be good reasons to accept a hypothesis even without conclusive evidence. This is especially so if the harms done by not doing anything are very high. Examples of such cases are Ebola *ca suffit!* and AZT trials. Obviously, there is a huge difference between these two cases: vaccination against Ebola eventually decreased harms, while AZT eventually increased them. This is a difference identified in hindsight. The discrepancy does not mean that reasons that lead to a decision to

approve these therapies for use are necessarily bad reasons, it rather means that decisions have to be assessed on a case-to-case basis taking a variety of factors into account, and that decision-making is particularly hard in biomedical contexts.

Such an interplay of epistemic and non-epistemic aims and goals, as well as of evidence and values, both epistemic and non-epistemic, is in line with Brown's pragmatist functionalism of inquiry and the rejection of lexical priority of evidence over values. Different values contribute to the fulfilment of epistemic aims because we have good reasons to care for them. Speed is one of them, together with non-epistemic and pragmatic considerations (understood as related to limitations of resources) that favor it. The value of speed does not point to true conclusions per se, but it rather contributes to truth in combination with an adequate degree of accuracy. In the next chapter I will focus more closely on the concept of time-sensitivity in science and the value of speed which addresses it.

4. Time-sensitivity in science

In this chapter I will examine the role of time-sensitivity in science by drawing on a discussion between Kevin Elliott and Daniel McKaughan (2014) and Daniel Steel (2016), on the role of non-epistemic values in theory assessment and the epistemic status of speed. Speed of generating results is a value instantiated in scientific practices, either through particular methods or through a decision to stop further testing. The focus in this chapter will be on methods. I will argue, first, that speed takes priority over ease of use in the cases which Elliott and McKaughan discuss; second, that speed is an epistemic value; and third, that Steel's account of values does not successfully distinguish extrinsically epistemic from non-epistemic values. Finally, I will propose an account of time-sensitivity and with the help of Biddle's (2013b) terminology apply it to case studies discussed in this thesis. An estimated degree of time-sensitivity should be understood as a contextual factor which bears on considerations about the limitations of temporal resources, and very directly informs the ways in which epistemic aims are pursued and achieved.

4.1 Introduction

Kevin Elliott and Daniel McKaughan (2014) argue that non-epistemic values sometimes legitimately take priority over epistemic values in assessing scientific theories, models, and

hypotheses because scientific representations are not only evaluated based on their fit with the world, but also based on the fit with the needs of their users. Their argument draws on accounts of scientific representation by Ronald Giere (2004) and Bas van Fraassen (2008), and two examples: expedited risk assessments of the toxicity of substances (Cranor 1993, 1995) and rapid assessment methods for wetland banking (Robertson 2004, 2006). Elliott and McKaughan consider speed and ease of use to be non-epistemic values. Their examples attempt to show that speed in the toxicity case and ease of use in the wetland banking case can have a more decisive role than that of being secondary considerations when epistemic values alone do not suffice to decide which representation to choose. They argue in favor of a direct (see Douglas 2000, 2009) or legitimate (see Steel 2010) role for non-epistemic values in the internal stages of scientific research and against the account of non-epistemic values as tie-breakers in resolving epistemic uncertainty, in other words, against the so called “epistemic priority approach”. Their criticism is different than Matthew Brown’s (2013) criticism of epistemic priority because what they propose as an alternative is “the aims approach”. Their aims approach is based on two guiding principles:

“First, if nonepistemic values are to play an appropriate role in assessing scientific representations, those engaged in the assessments need to be explicit about the goals of their assessments and the roles that nonepistemic values played in the assessments as a result. Second, nonepistemic values should receive priority only to the extent that they advance the goals associated with the assessments that are in play.” (Elliott and McKaughan 2014, 15)

In a comment on their paper, Daniel Steel (2016) argues that both of their examples fail to show that epistemic values have been overridden by non-epistemic values, but are rather cases in which non-epistemic values serve as secondary considerations for resolving epistemic uncertainty (i.e. as tie-breakers). According to Steel, the cases in question are not examples of accepting an epistemically inferior option because the argument rests on two problematic implicit premises: that it is epistemically better to wait for results generated by a more reliable method if one exists, and that it is bad from an epistemic perspective to select a simpler, less detailed model over one that is more complex and more detailed. In fact, in his (2010) article Steel uses Cranor’s (1993) analysis to argue that non-epistemic values can influence scientific inferences without compromising epistemic ends. The problem he identifies with Elliott and McKaughan’s account is that the first premise overlooks the epistemic costs of extended suspension of judgment and therefore “threatens to entail the absurd result that scientists should

never accept any claim” (Steel 2016, 610), while the second premise violates the principle of Ockham’s razor. Since there are many epistemic purposes to which hypotheses can be put, some of which can favor simplicity, there is nothing epistemically wrong with choosing a simpler option. Moreover, Steel characterizes both cases as illustrative of time-sensitivity:

“Both illustrate what I will call *time-sensitivity*, wherein it may be better for practical or social reasons to accept the results of a quicker-but-less-reliable method rather than wait for a slower-but-more-reliable-one. In both instances, there is a pressing interest to draw inferences in a timely manner: the protection of public health in the first and the economic interest of not unduly delaying construction projects in the second.” (Steel 2016, 609)

In order to examine the concept of time-sensitivity in both Elliott and McKaughan’s account and Steel’s, I start by arguing against Elliott and McKaughan’s view that the two tokens, speed and ease of use, independently of one another represent the same type, namely a non-epistemic value which sometimes takes priority over epistemic values in assessing scientific representations. Besides the problem of labeling speed and ease of use as non-epistemic, I claim that in both cases speed takes priority over simplicity and ease of use – the methods are simple and easy to use in order to be fast and enable fast (soon and many) applications. Both case studies are in fact primarily about speed, as already the titles of Elliott and McKaughan’s chapters reveal: *Expedited Risk Assessments* and *Rapid Assessment Methods*. After that I argue that speed is an epistemic value, contrary to Elliott and McKaughan and closer to Steel, but I part from the latter in that I do not think that the terminology of values suffices for explaining decision making in science.

I proceed by examining a way to account for time-sensitivity with the help of Steel’s conceptual framework. He offers a version of epistemic values which purports to argue in favor of maintaining the epistemic/non-epistemic distinction, as well as to be useful for delineating legitimate from illegitimate influence of non-epistemic values in research, namely by distinguishing between extrinsically and intrinsically epistemic values (Steel 2010). It seems to be consistent with Steel’s account to consider time-sensitivity an extrinsic epistemic value, since he argues for a broad understanding of epistemic values:

“Epistemic values can be manifested by things other than theories and hypotheses, such as methods, social practices, and community structures.” (2010, 19)

In this case, time-sensitivity might be a value manifested by social practices. However, I show that Steel's account of values does not prove to be helpful for handling the epistemic/non-epistemic controversy because it fails to distinguish between extrinsic epistemic values and non-epistemic values, especially when their influence on scientific research is legitimate, i.e. when they do not obstruct the attainment of truth.

I will claim that time-sensitivity is not captured well in either of the contrasting notions of value distinctions and argue that time-sensitivity is not a value of methods, but of problems to be solved in their particular contexts. We implicitly or explicitly assign a degree of time-sensitivity to problems in their specific contexts, which is a judgment about when we want or expect to have results from a particular instance of research, but it is neither a judgment exclusively external nor internal to science, but a requirement of efficiency which is both truth seeking and temporally constrained. I will eventually adopt Biddle's (2013b) terminology of contextual factors which can account for time-sensitivity in science. Finally, I will relate it to his argument from transient underdetermination.

4.2 Ease of use contributes to speed of generating results

The first example presented in Elliott and McKaughan's paper is based on Carl Cranor's analysis (1993, 1995) of different modeling approaches for assessing risks posed by toxic substances that are not pesticides or pharmaceuticals. In the United States the burden of proof is on the government to show that these products should be restricted or removed from the market and not on the manufacturers that produce them. Cranor (1995) analyzes trade-offs between different modeling approaches for assessing risks and concludes that social costs of relying on risk-assessment procedures which are rather accurate but slow are greater than of less accurate but quicker methodologies. This conclusion is based on the case of California Environmental Protection Agency (CEPA) which used an expedited risk assessment methodology in the early 1990s and was able to estimate carcinogenic potency of 200 chemicals in an 8 month period, while the traditional methodology was able to assess only 70 chemicals in 5 years, though with greater accuracy. The expedited procedure is called the linearized multistage default method (LMS) – it uses a carcinogenic potency data base, State of California data selection procedures and state-mandated default assumptions to facilitate otherwise time-consuming and science-intensive tasks in estimating dose-response relationships. Cranor calculates the difference between false positives and false negatives using different estimates,

some more and some less favorable to the expedited approach. It turns out to be a better approach in every case, in terms of minimizing social costs connected to underregulation of likely carcinogens. Elliott and McKaughan's conclusion is that speed is in this case prioritized over accuracy.

The second case deals with Rapid Assessment Methods (RAMs) for assessing similarity between different wetlands as part of mitigation measures when damaging or drying wetland areas. A destroyed wetland has to be compensated by preserving or restoring another wetland area, and regulatory agencies have to decide whether the destroyed and preserved wetlands are sufficiently similar so that the two could be traded. In recent years a mitigation "banking" system is developed by regulatory agencies, developers and entrepreneurs to handle mitigation. Geographer Morgan Robertson (2004, 2006) analyzes different methods to show how the banking method differs from the methods one would use if the goal was a detailed ecological characterization. Developers purchase mitigation "credits" from specialists who create "banks" of preserved or restored wetlands, in which they focus on specific features that are considered relevant for establishing the classification of "equivalence" between wetlands. RAMs consist of algorithms that convert data about a wetland into a numerical score that estimates a wetland's functional value and is typically represented by one main score rather than a variety of different scores "in order to keep the process simple." (Elliott and McKaughan 2014, 13) This case is supposed to be illustrative of ease of use as a value that is here taking priority over predictive accuracy. Their overall conclusion is that non-epistemic values sometimes take priority over the epistemic ones.

Against this, I argue that in these two cases, we are misled to judge speed and ease of use on a par with each other, as two tokens of the same type (a non-epistemic value that trumped predictive accuracy in assessing scientific representations), when in fact we have two cases of favoring an expedited outcome, which is prioritized over ease of use. What is going on is that ease of use is contributing to, and thus enabling, speed of generating results.⁷⁸ Speed is a value that has a decisive role of taking priority over predictive accuracy, if one wants to agree that this is what happens here, while ease of use and simplicity have only a transitive role as a means to achieve faster outcomes and applications. I do not want to imply that speed is always dependent on ease of use or that the benefits of ease of use and simplicity reduce to speed, but I claim that this is what is going on in the two examples by Elliott and McKaughan.

⁷⁸ While simplicity contributes to ease of use, but this relation is not of our interest at the moment.

For example, a theory can be simple and easy to use, but it can hardly be fast. It would be strange to claim that Euclidean geometry is faster than non-Euclidean geometry, or that Newtonian mechanics is faster than quantum mechanics. Although, as we could see from the discussion of gravitational waveform modeling, their usage in practice can contribute to a difference in speed of generating results. Thus post-Newtonian analytical approximation is relatively fast because it uses approximations and less computing power, as compared to numerical relativity which is computationally more intensive. Similarly, in the two cases discussed by Elliott and McKaughan we are not dealing with theories, but rather with methods and scientific practices that use simplifications, defaults, and idealizations designed to be applied to problems in particular contexts, and these methods and practices will most likely have simplicity and ease of use contributing to speed.

Elliott and McKaughan explicitly set out to show how non-epistemic values sometimes trump epistemic values such as predictive accuracy. Values that have supposedly done that in their examples are speed and ease of use. Although the second example is about making wetland models easy to use rather than making them highly accurate, the reason for making them easy to use is to make them readily available and thus faster to use. RAMs or “rapid assessment methods” are called rapid precisely because of that, but still the argument put forward is that ease of use is the value that took priority over accuracy in this case. It is certainly a feature of the RAM method in comparison to a more sophisticated one, but Elliott and McKaughan decided to talk about non-epistemic values in general based on the sample of two values which on the closer look turn out to be cases in which one value enables the other, i.e. ease of use enables speed, while it is itself facilitated by simplicity.

We can see the connection between simplicity, ease of use, and speed in both cases. Expedited risk assessment method is less science- and time-intensive, RAMs are easy to use because they are simple, and therefore results are generated faster than it would be the case with methods more detailed, complex, or difficult to handle. Methods do not generate results faster in order to be easier to use, but if they are easy to use, it can contribute to the speed of their performance. It is clear that being easy to use and being fast does not mean the same, but easy is in Elliott and McKaughan’s examples present as a value that enables speed, rather than the other way around, and also different from the cases in which they have no relation of dependency at all.

4.3 Speed as an epistemic value

The status of speed of inference is disputed in the discussion. Elliott and McKaughan claim that speed is a non-epistemic value:

“The cases discussed in the following sections focus on conflicts between the epistemic value of accurate prediction versus non-epistemic values such as ease of use or speed of generating results.” (2014, 7)

In his comment (2016), but also in an earlier article (2010), Steel argues that speed is an epistemic value:

“The trade-off between the speed and reliability of scientific methods, therefore, is a trade-off between two epistemic values.” (2010, 27)

First of all, not everything in science that we usually attribute values to can have the value of speed. It was already said that theories and hypotheses cannot be fast, but methods, and more broadly, practices, can. Methods, together with theories, models, hypotheses (representations) constitute practices in science, and practices can trade off speed and accuracy depending on their applications to problems in particular contexts. Speed, together with ease of use, is therefore a feature of methods and broader, a feature of practices as applied to problems in contexts. Problems need to be solved so the efficiency of methods and practices becomes important and has a bearing on the balance between values internal and external to the scientific practice that addresses them.

Steel’s distinction between epistemic “building blocks” and epistemic “endpoints” is useful here. Basic science is a building block for future research so it has a slower and more cautious approach when it comes to balancing reliability and speed of inference, because an error in that context is more likely to have damaging effects by leading to more errors. In contrast, scientific results that “are used primarily for some practical purpose, such as setting allowable exposure levels to toxic chemicals or predicting climate trends (...) are more like scientific endpoints than building blocks for future knowledge” (Steel 2010, 27).

Speed of scientific practices in generating results is an internal, epistemic value of scientific research – there is always a certain speed at which methods and practices operate. We might be tempted to call it non-epistemic because motivations to prioritize speed often come from outside of science and can operate at the expense of accuracy. But when speed is understood as speed of getting at true, or approximately true results, then it has a clearly epistemic role because it

moves us temporally closer to truth, i.e. it enables us to get in the possession of knowledge earlier and therefore advances our epistemic status. The non-epistemic part is still confined to different social reasons such as protection of health or economic benefits that instrumentalize speed for their ends at expense of accuracy. However, speed is often a means to promote these ends without epistemic costs, as Steel argues:

“To illustrate, suppose that the most epistemically reliable method will not produce any results for at least 1,000 years. Then a community that refuses to accept the results of less reliable methods may be at a distinct epistemic disadvantage because they remain ignorant while the knowledge of other communities that adopt less strict epistemic standards can advance by building on prior, albeit tentative, achievements.” (Steel 2016, 610)

Non-epistemic influence on scientific practices is not exerted by speed itself. Speed is a feature of a method or a practice and it belongs to the internal part of science all along the way. As such, it has to be traded off against other epistemic values. If non-epistemic reasons push the research in direction that moves it away from the truth, they can distort the balance between different values, for example prioritize speed of getting at any results over accuracy, which is problematic. But it can also happen that their influence on the trade-off is harmless and even beneficial, as I will explain in the next section. Social reasons are the non-epistemic part here, not the speed that they instrumentalize. Similarly to the biomedical research case, speed is not a problem per se, the problem is rather the quality of evidence acquired by a particular method or acquired by the time the decision was made to stop gathering further evidence.

4.4 Extrinsically epistemic is equal to non-epistemic-but-legitimately-influencing

Steel’s notion of epistemic values (2010) defines epistemic in terms of either intrinsically or extrinsically promoting the attainment of truth. Moreover, it allows that epistemic values are manifested by methods, social practices, and community structures. A value that Steel analyses at length as an example of extrinsic epistemic value is simplicity.

Simplicity is an extrinsic epistemic value for it can be truth-promoting, but only in combination with some other intrinsic epistemic value like accuracy, at least a sufficient degree of it. Extrinsically epistemic status saves its epistemic role without commitments to generality, because circumstances matter. In contrast, empirical accuracy is an intrinsic epistemic value,

and also a robust one, “in the sense of being epistemic in almost any setting”, while most other epistemic values Steel calls contextual because “their capacity to promote the attainment of truth depends on occurring within a specific set of circumstances” (2010, 20). Similar to simplicity, Steel would be consistent to argue that speed is a contextual and extrinsic epistemic value because it can promote the attainment of truth, but that depends on the appropriate degree of accuracy involved. In both cases discussed earlier it is precisely such a value, for it has an epistemic role granted by an accompanying degree of accuracy. This role consists in avoiding the cost of suspended judgment, which is avoiding a situation that does not bring us closer to truth.

Steel’s account of epistemic is not contrasted with evaluative, social, historical or contingent judgments/values (Steel 2010, 23), so it allows a broader scope of factors to count as extrinsically epistemic values, such as fundability or diversity of viewpoints, as long as they play a role in attaining the truth. This is why I contend that time-sensitivity might be considered as one of Steel’s extrinsically epistemic values. In the two cases from the beginning of the chapter, time-sensitivity was motivated by non-epistemic considerations, but since it did not compromise epistemic norms, even more, it promoted speed and therefore served an epistemic purpose of moving us temporally closer to truth, it can certainly be characterized as extrinsically epistemic by promoting the attainment of truth in the given circumstances.

The problem with Steel’s account is that it fails to discern between extrinsically epistemic values and non-epistemic values, especially when their influence is legitimate. The central aim of his account is to save the epistemic/non-epistemic distinction because of its usefulness in the argument from inductive risk. In order to do that, he develops “a principled basis for separating legitimate from illegitimate influences of non-epistemic values in scientific inference” (Steel 2010, 14), which states that “influences of non-epistemic values on scientific inferences are epistemically bad if and only if they impede or obstruct the attainment of truths.” (ibid., p. 15) In other words, influences of non-epistemic values are epistemically harmless if they do not impede or obstruct the truth. In fact, if they are not only harmless, but also beneficial in guiding us towards truth, as I claim they can be and often are, we can call them extrinsically epistemic. Let us take a closer look.

Steel analyzes two cases in which influence of non-epistemic values is welcome, to show how this is possible. The first case is precisely about speed – how long to wait or how much data to collect before accepting or rejecting a hypothesis, and the other is about judging some mistakes

worse than others. I will limit this analysis to the first type of cases. We have already seen that favoring speed, i.e. not waiting and not collecting additional, more detailed data, can be epistemically beneficial. I see no reason to regard this case as non-epistemic-but-legitimately-influencing, when it fits perfectly well under the scope of extrinsically epistemic values. If the default position of speed is for Steel extrinsically epistemic, as I contend it is, then what is non-epistemic, for example in the expedited assessment case, is the protection of human health as a value that motivates expedited risk assessments in the first place. If it does not obstruct the attainment of truth, but often promotes it (we cannot help people by pursuing untruthful and time-insensitive practices), why wouldn't we grant it an extrinsically epistemic status as well? There is no reason for separating the status of speed and the protection of human health in this particular case when the only criterion is their relation to the attainment of truth. After all, the circumstances matter. The protection of human health in these circumstances meets the condition of an extrinsically epistemic value. This becomes even clearer if we contrast it to fundability or diversity of viewpoints whose default position in Steel's account is extrinsically epistemic. There seems to be no problem in calling fundability and diversity of viewpoints non-epistemic-but-legitimately-influencing in some cases. There is no grounded difference between this status and an extrinsically epistemic status.

Steel's motivation is clear: he wants to save the argument from inductive risk which claims that non-epistemic values sometimes legitimately influence scientific research. More precisely, their influence is legitimate when deciding on which evidential standards to adopt and when to stop further testing. If non-epistemic consequences are grave, this awareness legitimately influences the choice to set higher standards for accepting a hypothesis. While I agree that this influence is legitimate, I also contend that in those cases we can understand non-epistemic values, such as protection of health, to be extrinsically epistemic. They do not impede or obstruct the attainment of truth and they often point in the direction of truth as, for example, in the case of time-sensitive practices, speed of getting at true results that they promote, and the protection of human health and economic benefits that motivate these time-sensitive practices.

Introducing the intrinsic/extrinsic distinction does not save the distinction between epistemic and non-epistemic values. Now there is no proper scope for non-epistemic-but-legitimately-influencing, because extrinsically epistemic values have appropriated it, along with some of the values that used to be encountered on the lists of epistemic values, like simplicity and external consistency. Either there are only intrinsic epistemic values (namely, only empirical accuracy and internal consistency), and everything else is sometimes extrinsically epistemic (when it

directs towards the truth in the given circumstances), otherwise it is non-epistemic because it does not have anything to do with the truth-seeking endeavor; or there are robust, intrinsic epistemic values and everything else is non-epistemic, but sometimes legitimately influencing scientific research. In any case, one side of the dichotomy has to be broadly construed, be it the epistemic or the non-epistemic side.

Steel endorsed a broad notion of epistemic values which does not fall in line with the usual epistemic side of the dichotomies (internal-external, fact-value, direct-indirect), but is constrained only by the relation to the attainment of truth. The alternative would be to be rigid on the epistemic side and count only intrinsic epistemic values as epistemic, and then carefully assess particular cases to allow for a legitimate influence of non-epistemic values in particular instances of research assessed on a case to case basis. Non-epistemic values would then have to be broadly construed to involve both simplicity and external consistency.

Steel's account implies a particularism about what is epistemic and what is non-epistemic in specific cases of scientific research. I do not think that this is bad news, but it does show that Steel's distinction between extrinsically and intrinsically epistemic does not deliver on its promises. Although a principled distinction between epistemic and non-epistemic values in terms of truth conduciveness as opposed to a lack of it can still hold, the application of Steel's account to particular instances of research is difficult because the respective domains of epistemic and non-epistemic change according to the circumstances of particular cases.⁷⁹ For example, promoting health will sometimes be extrinsically epistemic and sometimes non-epistemic, depending on the context. This works fine if we know the details of the cases, but most often we do not, and then the distinction does not help us because it is possible that non-epistemic-but-legitimately-influencing domain and extrinsically epistemic domain overlap. In this case, the inductive risk argument, which in Steel's account requires a clear distinction between epistemic and non-epistemic values, will not be satisfied, because we will not be able to say if non-epistemic values legitimately influence research or if they do not, since they will be indistinguishable from (extrinsically) epistemic values.

⁷⁹ I thank David Hopf for pointing out that particularism does not mean that the distinction does not hold, only that it is differently instantiated in particular cases. A parallel can be made with moral particularism. The fact that we have to assess moral questions on a case-to-case basis does not mean that we lack a criteria for establishing what is morally wrong or right in particular cases, only that we cannot generalize across different cases.

4.5 What is a value?

Values are, notably, a remarkably underdeveloped notion. Steel offers a clarification of his understanding of values:

“Values function as what Solomon calls “decision vectors,” that is, factors that “influence the outcome (direction) of a decision” (2001, 53). However, not all decision vectors would normally be thought of as values. For example, sexist bias is a decision vector but would not be a value in a social setting in which sexism is frowned on. Values, then, are decision vectors that are favorably regarded in a community.” (Steel 2010, 21)

Miriam Solomon’s “decision vectors” can be both values and preferences that influence a decision. She distinguishes between empirical decision vectors, which are “causes of preference for theories with empirical success, either success in general or one success in particular”, and non-empirical decision vectors, which are “other reasons or causes for choice” (Solomon 2001, 56). Both Steel’s and Solomon’s definitions are broad and acknowledge that a value in the context of decision making in science includes factors that are different from values narrowly construed. Boaz Miller gives another broad definition of value:

“A value is anything that serves as a basis for discriminating between different states of affairs and ranking some of them higher than others with respect to how much they are desired or cared about or how the personal, social, natural, or cosmic order ought to be.” (Miller 2014, 70)

Again, a value is “anything”: a preference, a characteristic, a feature of context, a means to grapple with a limitation of resource – that serves as a basis for discriminating between states of affairs. We have already seen that Brown’s (2013) account emphasizes reflective judgments about values and thus requires reasons for value judgments, which can be good or bad, and even empirically testable. Thus, values are both preferences and judgments, as well as various other factors, which is made explicit by Biddle (2013b). Biddle’s account is also against “the ideal of epistemic purity” or the value-free ideal. The reason for the terminological change is the confusion around the proper meaning of “value” and thus the scope of the science and values debate. His “contextual factors”, I argue, can capture the role of time-sensitivity better than “values”, regardless of how one defines a value.

“I have been discussing an ideal that I call the “ideal of epistemic purity”, rather than the “ideal of value-free science.” I have made this terminological change because there

are many contextual factors that might influence the epistemic appraisal of research, not just values. (...) Clearly, defenders of the “ideal of value-free science” – and defenders of almost any ideal for science, for that matter – wish to exclude factors such as having a hangover from influencing the appraisal of theories; and yet it is equally clear that having a hangover is not a value. Many critics of the “ideal of value-free science” wish to argue that ethical or political values can, in certain situations, legitimately influence theory or hypothesis appraisal; at the same time, most of these critics maintain that, in many situations, contextual factors that are not accurately described as values will invariably influence theory choice. None of this is to say that the question of the appropriate role of values in science isn’t important; it most definitely is. But the philosophical issues at the heart of the “science and values” debate extend beyond this question, and the terminology employed in this debate should reflect this fact.” (Biddle 2013b, 131)

Biddle makes a good point that the fact that a researcher has a hangover will inappropriately affect how she assess the evidence, but that a hangover cannot be called a value. Also, that the terminology employed in discussions of decision making in science should cover all influences on research, and not just value influences in terms of preferences and judgments. I believe that time-sensitivity is a good example of such a contextual factor that influences decision making and is not adequately described as a value. However, a high degree of time-sensitivity can give rise to a preference for a faster practice, be it through a preference for a faster method of generating results or by making a decision to stop further testing.

4.6 Time-sensitivity and transient underdetermination

The debate in which the notion of time-sensitivity is introduced provides us with understanding of both its non-epistemic setting and its epistemic directedness. Limitations of resources influence the choice of methods and practices. Sometimes we have social and pragmatic reasons to have the results quickly. Sometimes a scientist may want to have the results soon in order to move forward with her career or research, even if she is honestly dedicated to truth. Scientific work is embedded in time frames: of funding, career stages, a lifetime, a generation or of several generations. Whatever the reasons may be, we will assign a desired time frame for achieving certain ends in sight, even when it comes to “building block” science. We want to see some results at some time. The assigned value of the desired time frame, such as 5-10 year goal in

translational medicine or “the next decade” in gravitational wave physics, is the level of time-sensitivity, and it can and does affect how different values pertaining to research practices are balanced against each other, most obviously speed and accuracy or reliability of methods. The aim of attaining the truth does not only inform our methodological choices, it also motivates action that inevitably happens in time. It most certainly reflects an epistemic end, but it is also intertwined with all kinds of values and reasons. It would be misleading to call it a value, because it enters decision making as a judgment that has a say on how different values “hang” together. Even if the level of time-sensitivity is very low, it still is present.

Contrary to Steel who takes the normative stance of the inductive risk argument as the starting point in arguing for the influence of non-epistemic values in theory assessment, Biddle’s (2013b) starting point is the descriptive claim of the underdetermination thesis. He takes up Philip Kitcher’s (2001) distinction between transient, permanent, and global underdetermination to argue that transient underdetermination poses a significant challenge for the ideal of epistemic purity. The thesis of transient underdetermination, as already mentioned, states that some theories are underdetermined by logic and currently available evidence. The thesis of permanent underdetermination, on the other hand, states that at least some theories are underdetermined by logic and all possible evidence to which scientists will ever have access to, while the thesis of global underdetermination states that all theories are permanently underdetermined (Kitcher 2001, 30-31; Biddle 2013b, 125). Biddle makes the argument that the weak underdetermination thesis, the one of transient underdetermination, suffices to undermine the ideal of epistemic purity, either because it is impossible to screen out all contextual factors from transiently underdetermined research, or because screening them out would be done at a significant cost. The latter case is captured by the inductive risk argument, i.e. the normative claim about the role of values in theory appraisal.

“There are many areas of current, cutting edge science that are underdetermined, at least transiently, by evidence. In many of these areas, especially those relevant for public policy making, hypotheses must be evaluated quickly, before all of the evidence is in. Suppose, for example, that we need to determine whether a particular chemical used in pesticides is sufficiently safe or has acceptable environmental impact, or whether a drug that is currently on the market should be taken off the market. Suppose, furthermore, that the available evidence does not unambiguously determine which hypothesis should be accepted, or which decision should be made. In these cases, we do not have the luxury of waiting until all of the evidence is in, as postponing a decision could result in severe

environmental degradation, loss of life, or other adverse effects. In situations such as this, there is a gap between evidence and hypothesis choice, and this gap is inevitably filled by contextual factors. Thus, in these cases, the ideal of epistemic purity fails.” (Biddle 2013b, 126)

Time-sensitivity is closely related to transient underdetermination because it determines how quick the decision has to be made. The degree of time-sensitivity is always present, regardless of whether it is estimated very low or very high. Highly time-sensitive issues prioritize speed, which should be understood as an epistemic value because it improves our epistemic status. Inappropriate means of achieving speed of practices can make it seem that speed is a non-epistemic value, but the problem in these cases comes from inadequate evidence assessment.

For example, the fact that industry may want shorter trials is not a problem per se, the problem is insufficient evidence that comes with it. If, on the contrary, the lack of evidence that would normally be gathered in a longer time frame would be supplemented by a variety of evidence which would make the trial more predictive (at least for some patients) albeit shorter, then shorter research time might work. Speed is in tension with reliability, but if the conditions for accepting a hypothesis are epistemically sound, then speed is rather an asset than a downside. Such a case-to-case analysis is required in both Longino’s and Steel’s account of values in science, and also Brown’s pragmatic functionalism requires a lot of empirical evidence when determining whether evidence and values (which are backed up by reasons and thus subjectable to critical scrutiny) have been adequately put together to bring about the desired outcomes. Particularism is inevitable, but so is pragmatism, which is well served by an attention to speed of practices.

5. Time-sensitivity and the two case studies

The concept of time-sensitivity is a feature of problems in contexts and comes in degrees. It can be captured by Biddle’s “contextual factors” which influence decision making in science, both when it comes to resolving inductive risk cases (how long to gather evidence) and in other cases of transient underdetermination when the gap between theory and currently available evidence is not closed. For example, when deciding on the particular goals and methods of detecting a gravitational wave, which eventually settled on cross-correlating two identical devices.

In basic science like gravitational wave physics, it takes a lot of time, computational power and extremely sensitive instruments to handle all the uncertainties related to the end-in-sight. Not long ago, the end-in-sight was the detection of gravitational waves. Time-sensitivity might have been estimated low at the beginning, especially since there are no immediate applications of the research; for now it has “only” yielded the benefit of better understanding of the universe and matter. In this case it was reasonable to expect decades of research without a robust result. However, with time passing by, time-sensitivity of the detection attempts has grown nevertheless, because of huge cognitive and material investments which at some point require payoffs. Time-sensitivity motivates the choice of goals, new procedures for error estimates in waveform modeling, adding of computational power, and refinements of the instruments. Speeding up means coming up with new ways of getting the result, only in this context the tolerance for huge time-spans is higher. However, such tolerance is also exhaustive if there is no measurable advancement. This will be reflected in the shortage of funds, as well as in the shortage of researchers’ interest.

This particular research succeeded: gravitational waves were detected after more than fifty years of research. However, there are no guaranties that every research will be as successful as that, and it especially will not be the case that 50 years will be an acceptable time-span for every research practice. In comparison, recent efforts around translation in biomedical sciences are in part a reaction to the fact that the average time-span between discovery and implementation of therapeutic practices, which has been estimated 17 years (Morris et al. 2011), is considered way too long. Unsurprisingly, since the deliverances and applications of biomedical sciences are expected with much greater urgency than that of gravitational science. This judgment is so strong that it initiated a new model of biomedical research, namely translational science, dedicated to speeding up of the “bench to bedside” process. Time-sensitivity does have a saying on what the next step is and which values to prioritize in different research contexts.

Time-sensitivity is present in the context in which scientific research is done, in the uses it has, and problems it aims to address. The degree of time-sensitivity is implicitly or explicitly estimated and it has a bearing on the trade-offs between different values, such as speed and accuracy. As we have seen, simplicity and ease of use transitively address time-sensitivity by contributing to speed of methods and practices. For example, post-Newtonian approximation, a method that is simpler and easier to use, contributes to the speed of waveform modeling. The relation between simplicity, ease of use, and speed is present in translational medicine as well. Recall the need for “faster” and “clearer” translational practices (p. 82). The following quote

on translational epistemology explicitly links translational efforts with the need to simplify gene regulatory network (GRN) models and make them easier to use, in order to facilitate the cooperative work of scientists from different academic backgrounds:

“For translation, a critical issue is to form the conceptualization at the right level of abstraction. The model must be sufficiently complex to permit the translational problem to be formulated within it to a degree sufficient for the application at hand and it must be simple enough that the translational problem is not obscured by too much structure, the necessary parameters can be well enough estimated, and the optimization is mathematically and computationally tractable. The desire for simplicity drives much of the work of engineers: reduce (compress) the model to achieve tractability while at the same time keeping sufficient information so that the resulting solution, while suboptimal from the perspective of the full model, is still acceptable. One need only look at the great effort expended on image compression to recognize the importance of model reduction. A basic way of overcoming computational limitations when designing intervention strategies on GRNs is to reduce the GRN model. It is important for the success of the translational enterprise that there be tight interaction between the scientist and engineer when it comes to model complexity.” (Dougherty 2009, 107-108)

Similar to Elliott and McKaughan’s cases, simplicity of the model contributes to ease of use which ultimately contributes to speed of practices, especially if these practices have to be carried out by a community of researchers who have not been trained in the same way. In translational medicine, speed is an explicitly favored value and conjoined science-society means of achieving it have been systematically encouraged: from collecting more and diverse evidence, technological advancements in high-throughput screening and omics research, to changes in the design of trials, backed up by an increase in investments and establishing academic-industrial partnerships. However, valuing speed does not excuse impermissible practices when speed is achieved at the expense of adequate evidence assessment. Such problematic cases are best described as failures of judgment about how to increase speed (When to stop a trial? How to conduct an adaptive design study?), not as failures of judgment that speed should be valued. Since we cannot detect particular motives – has speeding up been prioritized because of a concern for patients’ benefit, an increase in income, epistemic efficiency, or most probably because of a combination of all of these factors, we will have to settle that speed cannot be dismissed as a mere non-epistemic influence. This influence, though

lurking in the background, does not compromise the fact that when two methods are empirically adequate, the faster one should be prioritized.

A method can generate results faster in comparison to another, and those results can be more or less accurate, but how much the setting of this activity is time-sensitive is a contextual and evaluative judgment that gives rise to concerns about efficiency and has a saying on how different epistemic values are balanced against each other in particular instances of research. Furthermore, it has a saying on whether different non-epistemic values should influence the research, for example, by stopping further testing in order to reach a conclusion about a problem that needs to be addressed urgently. Time-sensitivity should therefore be understood as a contextual factor assessed with the aim of attaining the truth but informed by the peculiarities of here and now. Highly time-sensitive issues favor expedited methods, in other words: higher the time-sensitivity, more valuable the speed. More generally, pragmatic considerations such as those about the limitations of time enter decision making in science without compromising epistemic ends. Thus, pragmatic is epistemic.

6. Conclusion

In part III I have first introduced the concept of epistemic, non-epistemic, and pragmatic values, as well as accounts that employ them, such as the value-free ideal, the inductive risk argument, and the underdetermination thesis. I have taken up the discussion between Elliott and McKaughan (2014) and Steel (2016) on time-sensitivity in science and the epistemic status of speed of practices, in order to argue for an epistemic role for speed and against Steel's (2010) conceptualization of values. I have introduced Biddle's (2013b) terminology which can account for time-sensitivity as a contextual factor that influences decision making in science. I have argued that time-sensitivity is a feature of problems that have to be solved in their particular contexts, and that it has an influence on how epistemic values are traded off against each other, but also that non-epistemic values influence the estimated degree of time-sensitivity.

A high degree of time-sensitivity can prioritize speed of practices at expense of other values. Most importantly, speed is still understood as an epistemic value because it promotes epistemic ends by addressing the limitations of resources, namely, temporal resources. Speed of practices is an epistemic value which can, in different cases, be either adequately or inadequately traded off against other epistemic values, motivated by both epistemic and non-epistemic reasons. This account is in line with the established erosion of the distinction between internal and external

to science as exemplified in the two case studies. Since the concept of values faces problems, I opted for the term “contextual factor” which can better capture the limitations of resources which are at the heart of this thesis. Limitations of resources require that epistemic aims are pragmatically informed, meaning – informed about the availability of time, money, material, and cognitive and computational power. Pragmatic considerations that address these limitations and contribute to achieving a desired speed of practices in generating results are a necessary dimension of the pursuit and fulfilment of epistemic goals.

Such an account is in line with Brown’s (2013) rejection of the lexical priority of evidence over values, Douglas’ (2000, 2009) argument from inductive risk, Biddle’s (2013b) argument from transient underdetermination, and Steel’s (2010) implicit particularism about values. The argument from transient underdetermination can account for cases of inductive risk, but also broaden the influence of values to a role larger than the one they are allowed to play in the argument from inductive risk. It can allow that decisions in science are made jointly by evidence and values. Values can in principle be distinguished to epistemic and non-epistemic depending on their relation to truth, but what exactly is truth conducive in particular instances of research should be assessed on a case-to-case basis. Such a pragmatist and particularistic account does not bear a commitment to a sharp distinction between epistemic and non-epistemic values in terms of a stable and uniform list of values, and it can account for considerations of limited resources in the pursuit of epistemic aims. Moreover, it does justice to the value of speed of practices, which has normally been understood as a part of the borderlands between epistemic and non-epistemic. I concede that its status should be understood as epistemic, even though its prioritization is subjectable to non-epistemic influences. This, however, is the common fate of epistemic values that should be embraced rather than avoided.

Conclusion

The starting point of this dissertation has been a recognition that time is a limited resource. This recognition motivated the hypothesis that scientific practice should and does reflect these constrained circumstances. Two case studies have been chosen with the aim to represent very different research practices, and their analyses have shown that considerations about time limits influence different stages of research in both cases.

The focus has been on speed of practices and their ability to generate results, either in the form of methods or in the form of decisions about which goals to set, as well as decisions about when to stop further testing. The value of speed has been central to this dissertation because of its role in addressing limited time frames in which scientific practices inevitably operate. Faster methods make those who use them epistemically better off in a unit of time. After a discussion of two case studies in the first two parts of the thesis, the third part provides an overview of the accounts in current debates about values and science which involve, explicitly or implicitly, a category of pragmatic values. This is a category that in some accounts is able to capture the value of speed (McMullin 1982). Other accounts have also been presented, in which speed is understood either as an epistemic (Steel 2010, 2016) or a non-epistemic value (Elliott and McKaughan 2014). This notable discrepancy shows that a focus on speed of methods and practices, and more generally a focus on epistemological and ethical aspects of time in scientific research, is a fruitful and underrepresented topic in philosophy of science, especially in science and values debates. I hope that this dissertation can at least partly fill in the gap.

I have argued for an epistemic role for speed by taking up a recent discussion and focusing on the role of time-sensitivity. Time-sensitivity is a contextual factor which comes in degrees and influences how epistemic values are balanced against each other. Non-epistemic values, on the other hand, influence the estimated degree of time-sensitivity. This account supports the view that both epistemic and non-epistemic values should have a role in decision making in science, and that their interaction is particularly necessary when addressing pragmatic issues such as how to cope with limited resources. Moreover, that such pragmatic considerations are an inevitable part of the pursuit and fulfillment of epistemic aims.

The first case exemplifies fundamental research without immediate societal applications, open-ended in terms of both the timeline and research goals. My analysis has shown that reflections

on time and speed nonetheless play a role in all stages of research, as well as on different levels of research organization. The second case exemplifies applied research with high social stakes, explicitly dedicated to acceleration of discovery and research, and motivated by non-epistemic reasons. Inter- and extra-scientific means of achieving acceleration were presented, and the ethical dimension present in the trade-off between different harms was emphasized. What is internal and what external to science has, however, often been hard to distinguish. The case studies provide an insight into the methods of gravitational wave modeling, problems in drug discovery practices, and features of different designs of clinical trials. My aim was to learn from scientific practices about what is important in decision making in science. Time and speed have turned out to be an invaluable resource for a philosophical discussion, both on the descriptive and the normative level. I hope that this dissertation is a valuable contribution on both of these levels.

Bibliography

- Abbott, B. P. et al. (LIGO Scientific Collaboration and Virgo Collaboration) (2016a), “Observation of Gravitational Waves from a Binary Black Hole Merger”, *Physical Review Letters* 116 (6) 61102.
- Abbott, B. P. et al. (LIGO Scientific Collaboration and Virgo Collaboration) (2016b), “Properties of the Binary Black Hole Merger GW150914”, *Physical Review Letters* 116 (24) 241102.
- Adam, M. (2005), “Integrating research and development: the emergence of rational drug design in the pharmaceutical industry”, *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences* 36(3): 513-537.
- Adam, M. (2011), “Multi-Level Complexities in Technological Development: Competing Strategies for Drug Discovery”, in M. Carrier, A. Nordmann (eds.) *Science in the Context of Application*. 67-83. Netherlands: Springer.
- Adam, M.; Carrier, M. and Wilholt, T. (2006), “How to serve the customer and still be truthful: methodological characteristics of applied research”, *Science and Public Policy* 33(6): 435-444.
- All for Mila, n.d. Retrieved on November 24, 2020 from http://lifeformila.com/index_eng.html
- American Chemical Society National Historic Chemical Landmarks: “Tagamet®: Discovery of Histamine H₂-receptor Antagonists”, n.d. Retrieved on November 24, 2020 from <http://www.acs.org/content/acs/en/education/whatischemistry/landmarks/cimetidinetagamet.html>
- Andersen, H. (2012), “Mechanisms: what are they evidence for in evidence-based medicine?” *Journal of Evaluation in Clinical Practice* 18(5): 992-999.
- Aronson, J. K. (2017a, March 24), “When I use a word . . . Translational research – early developments”, *British Medical Journal* blogs. Retrieved from <https://blogs.bmj.com/bmj/2017/03/24/jeffrey-aronson-when-i-use-a-word-translational-research-early-developments/>

- Aronson, J. K. (2017b, March 31), “Jeffrey Aronson: When I use a word... Translational research – an early model”, *British Medical Journal* blogs. Retrieved from <https://blogs.bmj.com/bmj/2017/03/31/when-i-use-a-word-translational-research-an-early-model/>
- Aronson, J. K. (2017c, April 7), “Jeffrey Aronson: When I use a word... Translational research – a further model”, *British Medical Journal* blogs. Retrieved from <https://blogs.bmj.com/bmj/2017/04/07/jeffrey-aronson-when-i-use-a-word-translational-research-a-further-model/>
- Aronson, J. K. (2017d, April 14), “Jeffrey Aronson: When I use a word... Translational research – a new operational model”, *British Medical Journal* blogs. Retrieved from <https://blogs.bmj.com/bmj/2017/04/14/jeffrey-aronson-when-i-use-a-word-translational-research-a-new-operational-model/>
- Aronson, J. K. (2017e, April 21), “Jeffrey Aronson: When I use a word... Defining translational research”, *British Medical Journal* blogs. Retrieved from <https://blogs.bmj.com/bmj/2017/04/21/jeffrey-aronson-when-i-use-a-word-defining-translational-research/>
- Aronson, J. K. (forthcoming), “The translation of pharmacological actions of medications into clinical outcomes” (Chapter 2.8). Author manuscript.
- Bacon, F. (2009 [1620]), *Novum organum and associated texts* (Reprinted). Rees, G. and Beal, P. (eds.) Oxford: Clarendon Press.
- Bagur, J. D. and Guissinger, A. S. (1987), “Technology transfer legislation: an overview”, *Journal of Technology Transfer* 12(1): 51-63.
- Bartfai, T. and Lees, G. (2013), *The Future of Drug Discovery. Who Decides Which Diseases to Treat?* Academic Press.
- Beckett, A. G. and Stevenson, C. J. (1956), “Unusual neurological complications after sudden withdrawal of cortisone”, *The British Medical Journal* 2(4983): 27-28.
- Berwick, D. M. (2005), “Broadening the view of evidence-based medicine”, *Quality and Safety in Health Care* 14(5): 315-316.
- Biddle, J. (2007), “Lessons from the Vioxx Debacle: What the Privatization of Science Can Teach Us About Social Epistemology”, *Social Epistemology* 21(1): 21-39.

- Biddle, J. (2013a), “Institutionalizing Dissent: A Proposal for an Adversarial System of Pharmaceutical Research”, *Kennedy Institute of Ethics Journal* 23(4): 325-353.
- Biddle, J. (2013b), “State of the field: Transient underdetermination and values in science”, *Studies in History and Philosophy of Science* 44: 124-33.
- Biddle, J. (2016), “Inductive Risk, Epistemic Risk, and Overdiagnosis of Disease”, *Perspectives on Science* 24(2): 192-206.
- “Bigger Than Life”, n.d. *Wikipedia*. Retrieved on November 24, 2020 from https://en.wikipedia.org/wiki/Bigger_Than_Life
- Bilkey et al. (2019), “Optimizing Precision Medicine for Public Health”, *Frontiers in Public Health* 7(42): 1-9.
- Boenink, M. (2010), “Molecular medicine and concepts of disease: the ethical value of a conceptual analysis of emerging biomedical technologies”, *Medical Health Care and Philosophy* 13: 11-23.
- Broach, J. R. and Thorner, J. (1996), “High-throughput screening for drug discovery”, *Nature* 384: 14-16.
- Brown, J. R. (2004), “Money, Method and Medical Research”, *Episteme* 1(1): 49-59.
- Brown, J. R. (2017), “Socializing Medical Research”, in K. Elliott and D. Steel (eds.) *Current Controversies in Values in Science*. 147-161. New York and London: Routledge.
- Brown, M. J. (2013), “Values in Science beyond Underdetermination and Inductive Risk”, *Philosophy of Science* 80(5): 829-839.
- Burns, C. M. (2016), “The History of Cortisone Discovery and Development”, *Rheumatic Disease Clinics of North America* 42(1): 1-14. doi:10.1016/j.rdc.2015.08.001.
- Bush, V. (1945), *Science, the Endless Frontier. A Report to the President*. Washington: United States Government Printing Office.
- Bueter, A. (2019), “Epistemic Injustice and Psychiatric Classification”, *Philosophy of Science* 86(5): 1064-1074.
- Calain, P. (2018), “The Ebola clinical trials: a precedent for research ethics in disasters”, *Journal of Medical Ethics* 44: 3-8.

- Canali, S. (2016, December), “Big Data, epistemology and causality: Knowledge in and knowledge out in EXPOsOMICS”, *Big Data & Society*.
doi:10.1177/2053951716669530.
- Carlisle, J. M. (1950), “Cortisone (compound E); summary of its clinical uses”, *British Medical Journal* 2(4679): 590-595.
- Carpenter, D.; Zucker, E. J. and Avorn, J. (2008), “Drug-Review Deadlines and Safety Problems”, *New England Journal of Medicine* 58: 1354-13561.
- Carrier, M. and Finzer, P. (2011), “Theory and Therapy: On the Conceptual Structure of Models in Medical Research”, in M. Carrier and A. Nordmann (eds.) *Science in the Context of Application*. 85-101. Netherlands: Springer.
- Cartwright, N. (2011), “The Art of Medicine. A philosopher’s view of the long road from RCTs to effectiveness”, *The Lancet* 377: 1400-1401.
- Cartwright, N. (2012), “Will This Policy Work for You? Predicting Effectiveness Better: How Philosophy Helps”, *Philosophy of Science* 79(5): 973-989.
- Castelvecchi, D. (2019, January 2), “Japan’s pioneering detector set to join hunt for gravitational waves”, *Nature* 565, 9-10. Retrieved from <https://www.nature.com/articles/d41586-018-07867-z>
- Clarke, B.; Gillies, D.; Illari, P.; Russo, F. and Williamson, J. (2013), “The evidence that evidence-based medicine omits”, *Preventive Medicine* 5(6): 745-747.
- Cohrs, R. J. et al. (2015), “Translational Medicine definition by the European Society for Translational Medicine”, *New Horizons in Translational Medicine* 2: 86-88.
- Collins, F. S.; Wilder, E. L. and Zerhouni, E. (2014), “NIH Roadmap/Common Fund at 10 years”, *Science* 345(6194): 274-276.
- Collins, H. (2004), *Gravity’s Shadow: The Search for Gravitational Waves*. Chicago and London: The University of Chicago Press.
- Collins, H. (2011, 2013), *Gravity’s Ghost and Big Dog: Scientific Discovery and Social Analysis in the Twenty-First Century*. The University of Chicago Press.
- Collins, H. (2017), *Gravity’s Kiss: The Detection of Gravitational Waves*. Cambridge Massachusetts: MIT Press.

- Contopoulos-Ioannidis, D. G.; Alexiou, G. A.; Gouvas, T. C. and Ioannidis, J. P. (2008), “Life Cycle of Translational Research for Medical Interventions”, *Science* 321(5894): 1298-1299.
- Contopoulos-Ioannidis, D. G.; Ntzani, E. and Ioannidis, J. P. (2003), “Translation of highly promising basic science research into clinical applications”, *American Journal of Medicine* 114(6): 477-484.
- Cooksey, D. (2006) *A review of UK health research funding*. London: Crown, HM Treasury. Retrieved from https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/228984/0118404881.pdf
- Copeland, S. (2015), *The Case of the Triggered Memory: Serendipitous Discovery and the Ethics of Clinical Research* (Doctoral dissertation). Dalhousie University.
- Copeland, S. (2019), “On serendipity in science: discovery at the intersection of chance and wisdom”, *Synthese* 196(6): 2385-2406.
- Copeman, W. S. C. et al. (1950), “A Study of Cortisone and Other Steroids in Rheumatoid Arthritis”, *The British Medical Journal* 2 (4684): 849–855.
- Copeman, W. S. C. et al. (1952), “Observations on Prolonged Cortisone Administration in Rheumatoid Arthritis”, *The British Medical Journal* 1(4755): 397-403.
- “Cortisone And A.C.T.H. In Britain” (1950, December 16), *The British Medical Journal* 2(4693): 1375. Retrieved from <http://www.jstor.org/stable/25359296>
- Cranor, C. (1993), *Regulating Toxic Substances*. Oxford: Oxford University Press.
- Cranor, C. (1995), “The Social Benefits of Expedited Risk Assessments”, *Risk Analysis* 15: 353-58.
- Crewe, T. (2018), “Here was a plague. Review of the books *How to Survive a Plague: The Story of How Activists and Scientists Tamed Aids*, *Patient Zero and the Making of the Aids Epidemic*, *Modern Nature: The Journals of Derek Jarman, 1989-90*, *Smiling in Slow Motion: The Journals of Derek Jarman, 1991-94* and *The Ward*”, *London Review of Books* 40(18): 7-16.

- CTSI Tufts Clinical and Translational Science Institute: “What is translational science?”, n.d. Retrieved on November 24, 2020 from <https://www.tuftsctsi.org/about-us/what-is-translational-science/>
- de Melo-Martín, I. and Intemann, K. (2016), “The Risk of Using Inductive Risk to Challenge the Value-Free Ideal”, *Philosophy of Science* 83(4): 500-520.
- de Melo-Martín, I. and Intemann, K. (2018), *The Fight Against Doubt: How to Bridge the Gap Between Scientists and the Public*. Oxford University Press.
- DiMasi, J. A.; Hansen, R. W. and Grabowski, H. G. (2003), “The price of innovation: new estimates of drug development costs”, *Journal of Health Economics* 22(2): 151-185.
- Dougherty, E. R. (2009), “Translational Science: Epistemology and the Investigative Process”, *Current Genomics* 10: 102-109.
- Douglas, H. (2000), “Inductive Risk and Values in Science”, *Philosophy of Science* 67(4): 559-579.
- Douglas, H. (2009), *Science, Policy, and the Value-Free Ideal*. Pittsburgh: University of Pittsburgh Press.
- Douglas, H. (2013), “The Value of Cognitive Values”, *Philosophy of Science* 80(5): 796-806.
- Douglas, H. (2017), “Why Inductive Risk Requires Values in Science”, in K. Elliott and D. Steel (eds.) *Current Controversies in Values in Science*. 81-94. New York and London: Routledge.
- EBM+, n.d. Retrieved on November 24, 2020 at <http://ebmplus.org/>
- Elliott, K. and McKaughan, D. (2014), “Nonepistemic Values and the Multiple Goals of Science”, *Philosophy of Science* 81(1): 1-21.
- Emery, C. (1949), “Cortisone”, *The British Medical Journal* 2(4628): 652-653.
- Epstein, S. (1996), *Impure Science. AIDS, Activism, and the Politics of Knowledge*. Berkeley: University of California Press.
- Fang, F. C. and Casadevall, A. (2010), “Lost in Translation – Basic Science in the Era of Translational Research”, *Infection and Immunity* 78(2): 563-566.

- FDA (2004), *Innovation/Stagnation: Challenge and Opportunity on the Critical Path to New Medical Products*. Retrieved from <https://www.who.int/intellectualproperty/documents/en/FDAproposals.pdf>
- FDA (2014, April 23), “FDA Drug Safety Communication”. Retrieved from <https://www.fda.gov/media/88483/download>
- FDA (2018a), *New Drug Therapy Approvals*. Retrieved from <https://www.fda.gov/files/about%20fda/published/2017-New-Drug-Therapy-Approvals-Report.pdf>
- FDA (2018b, July), “Biosimilar Product Information”. Retrieved from <https://www.fda.gov/drugs/biosimilars/biosimilar-product-information>
- FDA (2018c, April 1), “Step 3: Clinical Research”, Retrieved from <https://www.fda.gov/patients/drug-development-process/step-3-clinical-research>
- FDA (2018d, April 23), “Critical Path Initiative”. Retrieved from <https://www.fda.gov/science-research/science-and-research-special-topics/critical-path-initiative>
- FDA (2018e, October 31), “Adaptive Design Clinical Trials for Drugs and Biologics”. Retrieved from <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/adaptive-design-clinical-trials-drugs-and-biologics>
- Fishburn, C. S. (2013), “Translational research: the changing landscape of drug discovery”, *Drug Discovery Today* 18 (9-10): 487-494.
- Fogel, D. B. (2018), “Factors associated with clinical trials that fail and opportunities for improving the likelihood of success: A review”, *Contemporary Clinical Trials Communications* 11: 156-164.
- Fuller, J. (2016, April 25) “The Reading Room: ‘Making Medical Knowledge’”. *British Medical Journal* blogs. Retrieved from <https://blogs.bmj.com/medical-humanities/2016/04/25/the-reading-room-making-medical-knowledge/>
- Gabbard, H.; Williams, M.; Hayes, F. and Messenger, C. (2018), “Matching Matched Filtering with Deep Networks for Gravitational-Wave Astronomy”, *Physical Review Letters* 120 (14) 1103.

- Garralda et al. (2019), “New clinical trial designs in the era of precision medicine”, *Molecular Oncology* 13: 549-557.
- Gaudillière, J.-P. (2013), “From Propaganda to scientific marketing: Schering, cortisone, and the construction of drug markets”, *History and Technology* 29(2): 188-209.
- Geisbert, T. W. (2017), “First Ebola virus vaccine to protect human beings?” *Lancet* 389: 479-480.
- GEO600: “Advanced Technologies”, n.d. Retrieved on December 2, 2020 from <https://www.geo600.org/23441/advanced-technologies>
- GEO600 Gravitational-Wave Detector, n.d. Retrieved on December 2, 2020 from <http://www.geo600.org/>
- George, D. and Huerta, E. A. (2018), “Deep Learning for Real-time Gravitational Wave Detection and Parameter Estimation: Results with Advanced LIGO Data”, *Physics Letters B* 778: 64-70.
- Giere, R. (2004), “How Models Are Used to Represent Reality.” *Philosophy of Science* 71 (Proceedings): 742–752.
- Glyn, J. H. (1998), “The Discovery of Cortisone: A Personal Memory”, *British Medical Journal* 317 (7161): 822-823.
- Glyn, J. and Todd, J. (1954), “Cortisone And Aspirin In Rheumatoid Arthritis”, *The British Medical Journal* 1(4875): 1376-1377.
- Goldman, A. I. (1999), *Knowledge in a Social World*. Oxford: Oxford University Press.
- González, G. and Cavaglià, M. (2016, March), “LIGO Scientific Collaboration News”, *LIGO Magazine* 8: 4-5.
- Grant, B. (2017, November 30), “Philanthropic Funding Makes Waves in Basic Science”, *The Scientist*. Retrieved from <https://www.the-scientist.com/careers/philanthropic-funding-makes-waves-in-basic-science-30184>
- Harrington, J. L. (2011) *Translational Space: An Ethnographic Study of Stem cell Research* (Doctoral dissertation). University of Exeter, United Kingdom.

- Harrington, J. L. and Hauskeller, C. (2014), “Translational Research: An Imperative Shaping the Spaces in Biomedicine”, *Tecnoscienza: Italian Journal of Science & Technology Studies* 4(2): 73-91.
- Hasin, Y.; Seldin, M. and Lusis, A. (2017), “Multi-omics approaches to disease”, *Genome Biology* 18: 83.
- Hempel, C. G. (1965), “Science and Human Values”, in *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science*, by C. G. Hempel, 81-96. New York: Free Press.
- Henao-Restrepo, A. M.; Camacho, A.; Longini, I. M. et al. (2017), “Efficacy and effectiveness of an rVSV-vectored vaccine in preventing Ebola virus disease: final results from the Guinea ring vaccination, open-label, cluster-randomised trial (Ebola Ça Suffit!)” *Lancet* 389: 505–518.
- Henao-Restrepo, A. M.; Longini, I. M.; Egger, M. et al. (2015), “Efficacy and effectiveness of an rVSV-vectored vaccine expressing Ebola surface glycoprotein: interim results from the Guinea ring vaccination cluster-randomised trial”, *Lancet* 386: 857-866.
- Henao-Restrepo, A. M.; Longini, I. M.; Røttingen, J. A.; Kieny, M. P. and Edmunds, W. (2018), “Authors’ reply” to Metzger, G. W. and Vivas-Martínez, “Questionable efficacy of the rVSV-ZEBOV Ebola vaccine”, *Lancet* 391: 1021-1022.
- Hench, P. S. (1964), “The Reversibility of Certain Rheumatic and Non-Rheumatic Conditions by the Use of Cortisone or of the Pituitary Adrenocorticotrophic Hormone”, in *Nobel Lectures, Physiology or Medicine 1942-1962*. 311-341. Amsterdam: Elsevier Publishing Company.
- Hench, P. S.; Kendall, E. C.; Slocumb, C. H. and Polley, H. F. (1949), “The effect of a hormone of the adrenal cortex, cortisone (17-hydroxy-11-dehydrocorticosterone: compound E), and of pituitary adrenocorticotrophic hormone on rheumatoid arthritis and acute rheumatic fever”, *Transactions of the Association of American Physicians* 62: 64-68.
- Heusler, K. and Kalvoda, J. (1996), “Between basic and applied research: Ciba's involvement in steroids in the 1950s and 1960s”, *Steroids* 61(8): 492-503.

- Hill, A. (1954). Cortisone And Aspirin In Rheumatoid Arthritis. *The British Medical Journal*, 1(4876), 1437-1437.
- Hillier, S. G. (2007), “Diamonds are forever: the cortisone legacy”, *Journal of Endocrinology* 195(1): 1-6.
- Hobbs, G. (2019, March) “In the pulsar business: A new era for the Parkes radio telescope”, *LIGO Magazine* 14: 24-25.
- Holman, B. (2019), “Philosophers on drugs”, *Synthese* 196: 4363-4390.
- Horgan, R. P. and Kenny, L. C. (2011), “‘Omic’ technologies: genomics, transcriptomics, proteomics and metabolomics”, *The Obstetrician & Gynaecologist* 13: 189-195.
- Horton R. (2001), “Lotronex and the FDA: a fatal erosion of integrity”, *Lancet* 357: 1544-45.
- Howick, J. (2011), *The philosophy of evidence-based medicine*. Wiley-Blackwell, Oxford.
- Howick, J., Glasziou, P. and Aronson, J. K. (2009), “The evolution of evidence hierarchies: what can Bradford Hill's ‘guidelines for causation’ contribute?”, *Journal of the Royal Society of Medicine* 102(5): 186-194.
- Hoyningen-Huene, P. (2006), “Context of Discovery versus Context of Justification and Thomas Kuhn”, in J. Schickore and F. Steinle, (eds.) *Revisiting Discovery and Justification*. 119-131. Netherlands: Springer.
- IndIGO (2011), *LIGO – India: Proposal for an Interferometric Gravitational Wave Observatory*. Retrieved from https://dcc.ligo.org/public/0075/M1100296/002/LIGO-India_lw-v2.pdf
- “In Science Fields” (1950, February 25), *The Science News-Letter* 57(8): 120-121. Retrieved from <https://www.jstor.org/stable/3927554>
- “In Science Fields” (1951, March 17), *The Science News-Letter* 59(11): 168-169. Retrieved from <https://www.jstor.org/stable/3929131>
- Intemann, K. (2015), “Distinguishing between legitimate and illegitimate values in climate modeling”, *European Journal for Philosophy of Science* 5: 217-232.
- “intracrine.” (2007), *Dorland's Medical Dictionary for Health Consumers*. Retrieved from <https://medical-dictionary.thefreedictionary.com/intracrine>

- Isaacson, R. (2016, March), “The Transition of Gravitational Physics – From Small to Big Science. Part 2”, *LIGO Magazine* 8: 56-61.
- Jadreškić, D. (2016) “Some social aspects of discovery, synthesis and production of cortisone in the 1930s-1950s.”, *Acta medico-historica Adriatica* 14(2): 333-346.
- Jeffrey, R. C. (1956), “Valuation and Acceptance of Scientific Hypotheses”, *Philosophy of Science* 23(3): 237-246.
- Jena, A. B.; Calfee, J. E.; Mansley, E. C. and Philipson, T. J. (2009), “‘Me-Too’ Innovation in Pharmaceutical Markets”, *Forum for Health Economics and Policy* 12(1): 5-25.
- Jenkins, J. and Sampson, P. (1967), “Conversion Of Cortisone To Cortisol And Prednisone To Prednisolone”, *The British Medical Journal* 2(5546): 205-207.
- Jogalekar, A. (2012, November 26), “The perils of translational research”, *Scientific American*. Retrieved from <https://blogs.scientificamerican.com/the-curious-wavefunction/the-perils-of-translational-research/>
- Jolly, J. A. (1980), “The Stevenson-Wydler Technology Innovation Act of 1980”, *Journal of Technology Transfer* 5(1): 69-80.
- Johnstone D. (2006), “Translational science – a sexy title for pre-clinical and clinical pharmacology”, *British Pharmacological Society* 4:1-3, pA2 online. Retrieved from <http://www.pa2online.org/articles/article.jsp?volume=5&issue=&article=54>
- Jukola, S. (2019), “On the evidentiary standards for nutrition advice”, *Studies in History and Philosophy of Biology and Biomedical Sciences* 73: 1-9. doi: 10.1016/j.shpsc.2018.05.007.
- Jureško, G. (2019a, February 3), “Prava istina o Spinrazi: Kako je 'čudesni lijek' koji usporava strašnu bolest posvađao doktore i pacijente i postao političko, a ne medicinsko pitanje”, *Jutarnji list*. Retrieved from <https://www.jutarnji.hr/vijesti/hrvatska/prava-istina-o-spinrazi-kako-je-cudesni-lijek-koji-usporava-strasnu-bolest-posvadao-doktore-i-pacijente-i-postao-politicko-a-ne-medicinsko-pitanje/8335004/>
- Jureško, G. (2019b, July 12), “Odobrio HZZO: Lijek Spinraza prvih šest mjeseci bit će odobren svim oboljelima”, *Jutarnji list*. Retrieved from <https://www.jutarnji.hr/vijesti/hrvatska/video-odobrio-hzzo-lijek-spinraza-prvih-sest-mjeseci-bit-ce-odobren-svim-oboljelima/9114239/>

- Kagarise, M. J. and Sheldon, G. F. (2000), “Translational ethics: A perspective for the new millennium”, *Archives of Surgery* 135(1): 39-45.
- Kantor, L. W. (2008), “NIH Roadmap for Medical Research”, *Alcohol Research & Health* 31(1): 12-13.
- KBC Rijeka (2016), *Strategija kliničkog bolničkog centra Rijeka 2016-2020*. Retrieved from <http://kbc-rijeka.hr/wp-content/uploads/2016/04/Strategija-Klinickog-bolni%C4%8Dkog-centra-Rijeka.pdf>
- Kendall, E. C. (1964), “The development of cortisone as a therapeutic agent“, in *Nobel Lectures, Physiology or Medicine 1942-1962*. 270-288. Amsterdam: Elsevier Publishing Company.
- Kimmelman J. and London A. J. (2011), “Predicting Harms and Benefits in Translational Trials: Ethics, Evidence, and Uncertainty”, *PLoS Medicine* 8(3): e1001010.
- Kitcher, P. (2001), *Science, Truth, and Democracy*. New York: Oxford University Press.
- Kitcher, P. (2011), *Science in a Democratic Society*. Amherst, New York: Prometheus Books.
- Kline, S. J. and Rosenberg, N. (1986), “An Overview of Innovation”, in R. Landau and N. Rosenberg (eds.), *The Positive Sum Strategy: Harnessing Technology for Economic Growth*. 275-305. Washington, DC: National Academy Press.
- Knaapen, L. (2014), “Evidence-Based Medicine or Cookbook Medicine? Addressing Concerns over the Standardization of Care”, *Sociology Compass* 8(6): 823-836.
- Knispel, B. (2017, January 11) “One Dozen and One Neutron Stars – Einstein@Home discovers 13 new gamma-ray pulsars”, Einstein@Home. Retrieved from <https://einsteinathome.org/content/one-dozen-and-one-neutron-stars-%E2%80%93-einsteinhome-discovers-13-new-gamma-ray-pulsars#comment-156490>
- Kuhn, T. S. (1962), *The Structure of Scientific Revolutions*. The University of Chicago Press.
- Kuhn, T. S. (1977), “Objectivity, Value Judgment, and Theory Choice”, in *The Essential Tension: Selected Studies in Scientific Tradition and Change*. 320-339. University of Chicago Press.

- Kuhn, T. S. (1989), “Objectivity, Value Judgment, and Theory Choice”, in *Readings in the philosophy of science*, B. A. Brody and R. E. Grandy (eds.), 356-368. Englewood Cliffs, N.J.: Prentice Hall.
- Leonelli, S. (2016), *Data-Centric Biology: A Philosophical Study*. Chicago: The University of Chicago Press.
- Levi, I. (1960), “Must the Scientist Make Value Judgments?”, *Journal of Philosophy* 57(2): 345-357.
- LIGO (2016, February 11), “Gravitational Waves Detected 100 Years after Einstein’s Prediction”, News Release. 1-5. Retrieved from https://www.ligo.caltech.edu/system/media_files/binaries/310/original/LHO-NewsRelease-11Feb16-Final.pdf?1455201669
- LIGO (2018), *The LIGO Laboratory Charter (2019-2023)*. Retrieved from https://dcc.ligo.org/public/0002/M060323/013/M060323-charter_V13.pdf
- LIGO Caltech, n.d. Retrieved on November 25, 2020 from <https://www.ligo.caltech.edu/>
- LIGO: “Facts”, n.d. Retrieved on November 25, 2020 from <https://www.ligo.caltech.edu/page/facts>
- LIGO Magazine* (2016, March), “A Perfect Source: Timeline of GW150914”, 8: 8-14. Retrieved from <https://ligo.org/magazine/LIGO-magazine-issue-8-extended.pdf#page=8>
- LIGO Magazine* (2019, March), “The gravitational wave ‘forecast’ for Observing Run 3. A discussion”, 14: 6-11. Retrieved from <https://ligo.org/magazine/LIGO-magazine-issue14.pdf#page=6>
- LIGO: “Sources and Types of Gravitational Waves”, n.d. Retrieved on November 25, 2020 from <https://www.ligo.caltech.edu/page/gw-sources>
- LIGO: “Timeline”, n.d. Retrieved on November 25, 2020 from <https://www.ligo.caltech.edu/page/timeline>
- LIGO: “What is an Interferometer?” n.d. Retrieved on November 25, 2020 from <https://www.ligo.caltech.edu/page/what-is-interferometer>

- Linsay, P.; Saulson, P.; Weiss, R. and Whitcomb. S. (1983), *A Study of a Long Baseline Gravitational Wave Antenna System* (The Blue Book). Retrieved from https://dcc.ligo.org/public/0028/T830001/000/NSF_bluebook_1983.pdf
- Loewenberg, S. (2009), “The Bayh–Dole Act: A model for promoting research translation?” *Molecular Oncology* 3: 91-93.
- Longino, H. (1990), *Science as social knowledge*. Princeton: Princeton University Press.
- Longino, H. (1995), “Gender, Politics, and the Theoretical Virtues”, *Synthese* 104(3): 383-397.
- Longino, H. (1996), “Cognitive and Non-cognitive Values in Science: Rethinking the Dichotomy”, in L. Hankinson Nelson and J. Nelson (eds.), *Feminism, Science, and the Philosophy of Science*. 39-58. Dordrecht: Kluwer.
- Longino, H. (2002), *The Fate of Knowledge*. Princeton: Princeton University Press.
- Longino, H. (2004), “How Values Can Be Good for Science”, in P. K. Machamer and G. Wolters (eds.), *Science, Values, and Objectivity*. 127-142. University of Pittsburg Press.
- LSC (2018), *LIGO Scientific Collaboration Program 2018-2019*. Retrieved from https://dcc.ligo.org/public/0153/M1800085/007/LSCProgram2018_light.pdf
- LSC and Virgo (2017), “The Background Symphony of Gravitational Waves from Neutron Star and Black Hole Mergers”, 1-2. Retrieved from <https://www.ligo.org/science/Publication-GW170817Stochastic/flyer.pdf>
- LSC (LIGO Scientific Collaboration), n.d. Retrieved on November 25, 2020 from <https://www.ligo.org/>
- Lyons, L. (2013), “Discovering the Significance of 5σ ”, arXiv:1310.1284.
- Maienschein, J.; Sunderland, M.; Ankeny, R. A. and Scott Robert, J. (2008), “The Ethos and Ethics of Translational Research”, *The American Journal of Bioethics* 8(3): 43-51.
- Mason, H. L.; Meyers, C. S. and Kendall, E. C. (1936), “The chemistry of crystalline substances isolated from the suprarenal gland”, *Journal of Biological Chemistry* 114: 613–631.
- Mason, J. (2013), “Introduction of biosimilars: not to be confused with generics”, *Prescriber* 5 (October): 7-8.

- Max Planck Institute for Gravitational Physics (Albert Einstein Institute), n.d. Retrieved on November 25, 2020 from <https://www.aei.mpg.de/>
- Max Planck Institute for Gravitational Physics (Albert Einstein Institute) (2017, October 16), “First observation of gravitational waves from merging neutron stars: Researchers’ Voices.” Retrieved from <https://www.aei.mpg.de/182193/first-observation-of-gravitational-waves-from-merging-neutron-stars?page=2>
- Meadors, G. D. (2014), *Directed searches for continuous gravitational waves from spinning neutron stars in binary systems* (Doctoral dissertation). The University of Michigan.
- Mervis, J. (1991, November 24), “Funding of two science labs receives pork barrel vs peer review debate”, *The Scientist*. Retrieved from <https://www.the-scientist.com/news/funding-of-two-science-labs-revives-pork-barrel-vs-peer-review-debate-60355>
- Metzger, G. W. and Vivas-Martínez, S. (2018), “Questionable efficacy of the rVSV-ZEBOV Ebola vaccine”, *Lancet* 391: 1021.
- McMullin, E. (1982), “Values in Science”, *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association* 4: 3-28.
- Miller, B. (2014), “Catching the WAVE: The Weight-Adjusting Account of Values and Evidence”, *Studies in History and Philosophy of Science Part A* 47: 69-80.
- “Modifications Of Cortisone” (1955, June 25), *The British Medical Journal*, 1(4929): 1520. Retrieved from <http://www.jstor.org/stable/20364025>
- Molina, B. S. et al. (2009), “The MTA at 8 years: Prospective follow-up of children treated for combined-type ADHD in a multisite study”, *Journal of the American Academy of Child & Adolescent Psychiatry* 48(5): 484-500.
- Morris, Z. S.; Wooding, S. and Grant, J. (2011), “The answer is 17 years, what is the question: understanding time lags in translational research”, *Journal of the Royal Society of Medicine* 104: 510-520.
- Mullard, A. (2019), “2018 FDA drug approvals”, *Nature Reviews Drug Discovery* 18: 85-89.
- Munos, B. (2009), “Lessons from 60 years of pharmaceutical innovation”, *Nature Reviews Drug Discovery* 8: 959-968.

- Muscular Dystrophy UK: “Spinraza”, n.d. Retrieved on November 25, 2020 from <https://www.muscular dystrophyuk.org/campaign-for-independent-living/fast-track/spinraza/>
- NASA (2015, December 3), “LISA Pathfinder Carries Advanced NASA Thruster Tech”. Retrieved from <https://www.jpl.nasa.gov/news/news.php?feature=4779>
- “New Sources Of Cortisone” (1950, June 24), *The British Medical Journal* 1(4668): 1477-1478. Retrieved from <http://www.jstor.org/stable/25357472>
- “News and Notes” (1949, August 5), *Science* 110 (2849): 151-154. Retrieved from <http://www.jstor.org/stable/1676546>
- NIH (2014), *A Decade of Discovery: The NIH Roadmap and Common Fund (2004-2014)*. Retrieved from <https://commonfund.nih.gov/sites/default/files/ADecadeofDiscoveryNIHRoadmapCF.pdf>
- Noell, G.; Faner, R. and Agusti, A. (2018), “From systems biology to P4 medicine: applications in respiratory medicine”, *European Respiratory Review* 27: 170110.
- NSF (2017), “LIGO Fact Sheet: NSF and the Laser Interferometer Gravitational-Wave Observatory”. Retrieved from https://www.nsf.gov/news/special_reports/ligoevent/pdfs/LIGO_factsheet_2017_v01.pdf
- Ohme, F. (2012), *Bridging the Gap between Post-Newtonian Theory and Numerical Relativity in Gravitational-Wave Data Analysis* (Doctoral dissertation). Universität Potsdam.
- Owen, R.; Macnaghten, P. and Stilgoe, J. (2012), “Responsible research and innovation: From science in society to science for society, with society”, *Science and Public Policy* 39: 751-760.
- Pammolli, F.; Magazzini, L. and Riccaboni, M. (2011), “The productivity crisis in pharmaceutical R&D”, *Nature Reviews Drug Discovery* 10: 428-438.
- “Paths To Cortisone” (1951, August 18), *The British Medical Journal* 2(4728): 406-408. Retrieved from <http://www.jstor.org/stable/25376093>
- Pielke, R. A. Jr. (2007), *The Honest Broker: Making Sense of Science in Policy and Politics*. Cambridge University Press.

- Pitkin, M. and Sun, L. (2017, August), “The search for Continuous Waves”, *LIGO Magazine* 11: 10-12.
- Quirke, V. (2005), “Making British cortisone: Glaxo and the development of corticosteroids in Britain in the 1950s-1960s”, *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences* 36: 645-674.
- Reichenbach, H. (1938), *Experience and Prediction. An Analysis of the Foundation and Structure of Knowledge*. Chicago, Illinois: University of Chicago Press.
- Reichstein, T. (1936), “Über Bestandteile der Nebennieren-Rinde, VI. Trennungsmethoden, sowie Isolierung der Substanzen F. a. H und J”, *Helvetica Chimica Acta* 19: 1107-1126.
- Reiss, J. (2017), “Meanwhile, Why Not Biomedical Capitalism?”, in K. Elliott and D. Steel (eds.) *Current Controversies in Values in Science*. 161-176. New York and London: Routledge.
- Research America (2017), *U.S. Investments in Medical and Health Research and Development 2013-2016. A report outlining U.S. medical and health R&D expenditures*. Retrieved from https://www.researchamerica.org/sites/default/files/RA-2017_InvestmentReport.pdf
- Robertson, M. (2004), “The Neoliberalization of Ecosystem Services: Wetland Mitigation Banking and Problems in Environmental Governance”, *Geoforum* 35: 361-373.
- Robertson, M. (2006), “The Nature That Capital Can See: Science, State, and Market in the Commodification of Ecosystem Services”, *Environment and Planning D, Society and Space* 24: 367-387.
- Robinson, M. D. (2019), “Financializing epistemic norms in contemporary biomedical innovation”, *Synthese* 196: 4391-4407.
- Rolin, K. (2017), “Can Social Diversity Be Best Incorporated into Science by Adopting the Social Value Management Ideal?”, in K. Elliott and D. Steel (eds.) *Current Controversies in Values in Science*. 113-130. New York and London: Routledge.
- Rooney, Phyllis (2017), “The Borderlands between Epistemic and Non-epistemic Values”, in K. Elliott and D. Steel (eds.) *Current Controversies in Values in Science*. 31-45. New York and London: Routledge.

- Rudner, R. (1953), "The Scientist Qua Scientist Makes Value Judgments", *Philosophy of Science* 20(1): 1-6.
- Russo, F., and Williamson, J. (2007), "Interpreting causality in the health sciences", *International Studies in the Philosophy of Science* 21: 157-170.
- Sackett, D. et al. (1996), "Evidence based medicine: what it is and what it isn't", *The British Medical Journal* 312: 71-72.
- Saenger, A. K. (2010), "Discovery of the Wonder Drug: From Cows to Cortisone", *Clinical Chemistry* 56(8): 1349-1350.
- Sánchez-Serrano, I. (2006), "Success in translational research: lessons from the development of bortezomib", *Nature Reviews Drug Discovery* 5: 107-114.
- Sarett, L. H. (1946), "Partial synthesis of pregnene-4-triol-17(b), 20(b), 21-dione-3,11 and pregnene-4-diol-17(b), 21-trione-3, 11, 20 monoacetate", *Journal of Biological Chemistry* 162: 601-632.
- Sarewitz, D. (1996), *Frontiers of Illusion. Science, Technology, and the Politics of Progress*. Philadelphia: Temple University Press.
- Schachter, H. M.; Pham, B.; King J.; Langford, S. and Moher, D. (2001), "How efficacious and safe is short-acting methylphenidate for the treatment of attention-deficit disorder in children and adolescents? A meta-analysis." *Canadian Medical Association Journal* 165(11): 1475-1488.
- Sismondo, S. (2004), "Pharmaceutical Maneuvers", *Social Studies of Science* 34(2): 149-159.
- Slater, L. B. (2000), "Industry and Academy: The Synthesis of Steroids", *Historical Studies in the Physical and Biological Sciences* 30(2): 443-480.
- Simoni, R. D.; Hill, R. L. and Vaughan, M. (2002), "The Isolation of Thyroxine and Cortisone: the Work of Edward C. Kendall", *The Journal of Biological Chemistry* 277(21): 10.
- Sofaer, N. and Eyal, N. (2010), "The diverse ethics of translational research", *The American Journal of Bioethics* 10(8): 19-30.

- Sokal, A. M. and Gerstenblith, B. A. (2010), “The Hatch-Waxman Act: Encouraging Innovation and Generic Drug Competition”, *Current Topics in Medicinal Chemistry* 10(18): 1950.
- Solomon, M. (2001), *Social Empiricism*. Cambridge, MA: MIT Press.
- Solomon, M. (2011), “Just a paradigm: evidence-based medicine in epistemological context”, *European Journal for Philosophy of Science* 1: 451-466.
- Solomon, M. (2015), *Making Medical Knowledge*. Oxford University Press.
- Staley, K. (2017), “Decisions, Decisions: Inductive Risk and the Higgs Boson”, in K. Elliott and T. Richards (eds.) *Exploring Inductive Risk. Case Studies of Values in Science*. 37-55. Oxford University Press.
- Steel, D. (2010), “Epistemic Values and the Argument from Inductive Risk”, *Philosophy of Science* 77: 14-34.
- Steel, D. (2016), “Accepting an Epistemically Inferior Alternative? A Comment on Elliott and McKaughan”, *Philosophy of Science* 83: 606-612.
- Steele, K. (2012), “The Scientist qua Policy Advisor Makes Value Judgments”, *Philosophy of Science* 79(5): 893-904.
- Stegenga, J. (2015a), “Effectiveness of Medical Interventions”, *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences* 54: 34-44.
- Stegenga, J. (2015b), “Measuring effectiveness”, *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences* 54: 62-71.
- Stegenga, J. (2017), “Drug Regulation and the Inductive Risk Calculus”, in K. Elliott and T. Richards (eds.) *Exploring Inductive Risk. Case Studies of Values in Science*. 17-37. Oxford University Press.
- Stegenga, J. (2018), *Medical Nihilism*. Oxford University Press.
- Strimbu, K. and Tavel, J. (2010), “What are Biomarkers?”, *Current Opinion in HIV and AIDS* 5(6): 463-466.

- Takebe, T., Imai, R. and Ono, S. (2018), “The Current Status of Drug Discovery and Development as Originated in United States Academia: The Influence of Industrial and Academic Collaboration on Drug Discovery and Development”, *Clinical and Translational Science* 11(6): 597-606.
- Taylor, D. (2016), “The Pharmaceutical Industry and the Future of Drug Development”, in R. E. Hester and R. M. Harrison (eds.) *Issues in Environmental Science and Technology 41: Pharmaceuticals in the Environment*. 1-33. The Royal Society of Chemistry.
- The White House, Office of the Press Secretary (2015, January 30), “Fact Sheet: President Obama’s Precision Medicine Initiative”. Retrieved from <https://obamawhitehouse.archives.gov/the-press-office/2015/01/30/fact-sheet-president-obama-s-precision-medicine-initiative>
- Thorlund, K.; Haggstrom, J.; Park, J. J. H. and Mills, E. J. (2018), “Key design considerations for adaptive clinical trials: a primer for clinicians”, *British Medical Journal* 360: k698.
- Thorn, G. (1951), “The Development and Application of ACTH and Cortisone”, *Bulletin of the American Academy of Arts and Sciences* 4(5): 2-4.
- Thorne, K. (1992), “Sources of Gravitational Waves and Prospects for Their Detection”, in A. Janis and J. Porter (eds.) *Recent Advances in General Relativity. Essays in Honor of Ted Newman*. 196-230. Birkhäuser: Boston, Basel, Berlin.
- Tonelli, M. R. (1998), “The philosophical limits of evidence-based medicine”, *Academic Medicine* 73(12): 1234-1240.
- “Treatment With Cortisone” (1951, July 28), *The British Medical Journal* 2(4725): 222-223. Retrieved from <http://www.jstor.org/stable/25375893>
- Trøst Jørgensen, J. (2008), “From blockbuster medicine to personalized medicine”, *Personalized Medicine* 5(1): 55-63.
- van der Laan, A. L. and Boenink, M. (2015), “Beyond Bench and Bedside: Disentangling the Concept of Translational Research”, *Health care analysis: journal of health philosophy and policy* 23: 32-49.
- van der Scheer, L.; Garcia, E.; van der Laan; van der Burg, S. and Boenink, M. (2017), “The Benefits of Patient Involvement for Translational Research”, *Journal of health philosophy and policy* 25(3): 225-241.

- van Fraassen, B. (2008) *Scientific Representation: Paradoxes of Perspective*. Oxford: Oxford University Press.
- Varghese, J. (2018), “Influence and prioritization of non-epistemic values in clinical trial designs: a study of Ebola ça Suffit trial”, *Synthese*: 1-17. doi:10.1007/s11229-018-01912-0.
- Vogenberg, F. R., Barash, C. I. and Pursel, M. (2010), “Personalized Medicine. Part 1: Evolution and Development into Theranostics”, *Pharmacy and Therapeutics* 35(10): 560-576.
- Wehling, M. (2006), “Translational medicine: can it really facilitate the transition of research ‘from bench to bedside’?” *European Journal of Clinical Pharmacology* 62: 91-95.
- Wehling, M. (2008), “Translational medicine: science or wishful thinking?” *Journal of Translational Medicine* 6: 31.
- WHO (2015), “Ring Vaccination Trial Consortium Ebola ça Suffit! Ebola vaccine Phase 3 trial Guinea”. Retrieved from http://www.who.int/immunization/research/meetings_workshops/3_Ebola_Ring_vaccination_Phase_3_trial_2015.pdf
- WHO (2018, May 30), “Ebola vaccine provides protection and hope for high-risk communities in the Democratic Republic of the Congo”. Retrieved from <https://www.who.int/news-room/feature-stories/detail/ebola-vaccine-provides-protection-and-hope-for-high-risk-communities-in-the-democratic-republic-of-the-congo>
- Wilholt, T. (2006), “Design Rules: Industrial Research and Epistemic Merit”, *Philosophy of Science* 73: 66-89.
- Wilholt, T. (2009), “Bias and values in scientific research”, *Studies in History and Philosophy of Science* 40: 92–101.
- Williams, D. A. (2005), “The NIH Roadmap: Timing is Everything”, *Molecular Therapy* 11(2): 173.
- Wilson, D. (1950), “ACTH and Cortisone in Clinical Practice”, *The American Journal of Nursing* 50(10): 649-652.

Woodward, R. B.; Sondheimer, F. and Taub, D. (1951), "The total synthesis of cortisone", *Journal of the American Chemical Society* 73: 4057.

Wu, X. P. (2017), "Listening to the Universe with gravitational waves", *National Science Review* 4: 680.

Zhang, Y.; Diao, T. and Wang, L. (2014), "Quantitative Evaluation of Translational Medicine Based on Scientometric Analysis and Information Extraction", *Clinical and Translational Science* 7(6): 465-469.