Evaluating Neural Futures:
Good Technoscience and the Challenge of Co-Production

Matthew Sample

A dissertation
submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

University of Washington

2016

Reading Committee:
Alison Wylie, Chair
Sara Goering
Sheila Jasanoff
Andrea Woody

Program Authorized to Offer Degree:
Philosophy

University of Washington

Abstract

Evaluating Neural Futures: Good Technoscience and the Challenge of Co-Production

Matthew Sample

Chair of the Supervisory Committee:
Professor Alison Wylie
Department of Philosophy

If we look beyond just the hypotheses, models, or evidence of technoscience, there are a variety of entangled, normative issues to be examined. Science and engineering enable the creation of new identities, change existing ways of life, and reflect collective visions for society. Accordingly, I use this dissertation to suggest how philosophy of science can address this challenge, taking the "co-production" of knowledge and social order (Jasanoff 2004) as my starting point. I argue, first, that constructivist science and technology studies, rather than precluding philosophy, lay the foundation for ethically and politically-sensitive philosophy of science. Second, I assess promising theoretical frameworks from Helen Longino, Lorraine Code, and Heather Douglas; each provides resources to evaluate technoscience, but require some changes to avoid traditional philosophical blindspots. Third, I shift to a more detailed consideration of neural engineering, as a test case for my interdisciplinary methodology. Ultimately, I propose a pragmatist conception of "good" (rather than true) technoscience, adopt a modest understanding of scholarly expertise, and call for a new philosophy of the field.

Acknowledgements

Intellectual work is inevitably social, and an honest scholar sees the autobiography behind theory. The chapters that follow reflect important relationships and activities in my life over the last five years. The project first took shape when Alison Wylie challenged me to think about good science and gave me the theoretical frameworks to connect that ideal to real-world practice. The challenge became concrete when Sara Goering introduced me to neuroethics at the Center for Sensorimotor Neural Engineering. I'd like to thank these philosophers, my dissertation committee, and especially my committee chair, for their continual support and encouragement.

I also owe thanks to the resolutely philosophical community at the Harvard Program on Science, Technology, and Society, which was my scholarly home for two years. The STS fellows and their friendly but persistent calls for accountability gave me a new perspective on my own expertise as a philosopher. Sheila Jasanoff, as an author and a mentor, helped me to see how a empirically-nuanced approach to science and technology doesn't preclude a focus on truth, rightness, or justice.

If a dissertation committee can embody a discipline, or if a bibliography reveals an invisible college, then this document hints at a productive space that extends beyond philosophy. I hope, then, that it can be read as a transcontinental testament to the many people who helped me articulate my thoughts, and to the collective concerns and affinities that brought us together.

Contents

Introduction

From neuroscience to neuroprosthetic: scientific literature as the philosopher's canary

In 1969, a young Eb Fetz published a short paper in *Science*, "Operant Conditioning of Cortical Unit Activity." Over the next three decades, it would become the textual touchstone for the broad and well-funded technoscientific field of neural engineering, for brain-computer interfaces, for brain-to-brain control, and robotic exoskeletons. In that key paper, Fetz reports that if you provide a Rhesus monkey with an incentive of banana-flavored food pellets—this is not unusual for studying animal behavior—and a way to hear the activity of a single neuron in its *own* motor cortex—a system of audible clicks correlated to neural firing—then the animal can be trained to voluntarily activate that single neuron. Whenever the monkey happened to push the firing rate to a maximum, Fetz would give it a pellet, reinforcing that state of firing. Regardless of what the cortical unit was previously controlling within the brain, it could be repurposed by the monkey to control an external meter and win treats.

This is a curious sort of conclusion, I propose, because it reports the effects of a very unusual set of conditions. Primates, humans included, almost never have access to the firing rate a single neuron in their brains. Why would Fetz expend time and costly animal resources to investigate such an unusual state of affairs? The answer, it turns out, is not so exotic in context. Fetz writes that "the technique of conditioning the activity of specific central structures by direct operant reinforcement will be useful for investigating neural mechanisms." As one might expect of a neuroscience piece submitted to *Science*, Fetz aimed for better understanding of existing causal mechanisms. So it's no surprise that the only explicit argumentation in the paper takes the form of ruling out alternative explanations for Fetz' claimed causal mechanism, successful operant conditioning.

I draw the reader's attention to this paper to make a simple point. You might excuse a philosopher of science for interpreting this paper from a purely epistemological perspective, from the eyes of a causally-oriented neuroscientist. There is very little in the paper that should shock the philosopher out of his state of epistemological myopia, to distract him from comparing Fetz' reasoning to Mill's methods or to Peter Lipton's contrastive theory of explanation. There is no gesturing towards technological applications of neuroscientific knowledge or acknowledgments of military defense grants. Just one author, a monkey, and a goal of causal understanding via good explanations. Careful scholars of science and technology studies (STS) will reject such a narrow characterization of Fetz' paper, and appropriately so, but the text itself gives us very little with which to work. Reaching beyond a narrowly epistemological understanding of Fetz' claims would require additional resources: historical context, personal interviews, knowledge of relevant institutional structures, and so on. My task here is not to craft such a robust STS story, but to make a comparison.

If we jump forward a few decades, to the mid-2000s, scientific work on the neuroscience of Rhesus monkeys is published in an entirely different register. It is an explicitly value-laden and political register that cannot excuse a narrowly epistemological focus. Consider "Cortical control of a prosthetic arm for self-feeding," published in *Nature* in 2008. There, Velliste et al. claim "we show how monkeys (*Macaca mulatta*) use their motor cortical activity to control a mechanized arm replica in a self-feeding task." Like Fetz, these authors are not testing a preformulated hypothesis, but are making a claim of artifice. And their methodology is remarkably similar to that of Fetz; they use operant conditioning to train monkeys to repurpose a cortical neuron, although this time, the neuron is connected to a robotic arm rather than a simple feedback signal. The similarities end there, however.

Immediately in the abstract, the authors claim that a brain-machine interface "could be applied to restore arm and hand function for amputees and paralyzed persons." Their study is meant to test whether Fetz' technique can enable real-time control of a robotic limb. In the process, it presupposes a vision of society in which the negative disabilities associated with paralysis or limb loss are fixed by means of technological interventions – brain surgery! – on the individual. The conclusion, correspondingly, does not include the words "understanding" or "knowledge" or "mechanism." Rather, "this demonstration of multi-degree-of-freedom embodied prosthetic control paves the way towards the development of dexterous prosthetic devices that could ultimately achieve arm and hand function at a near-natural level." The reasoning behind this particular instrumental strategy is never made explicit, just as the criteria used to define "successful" or "near-natural" use of the robotic arm are left unanalyzed. One might find these presuppositions in need of some justification, especially considering the authors' claim that their findings with Rhesus monkeys "can be incorporated into future designs of prostheses for dexterous movement," for humans.

If these ethically and politically-loaded claims aren't enough to problematize a solely epistemological perspective on the paper, consider also that the multiple authors include a biomedical engineer, an electrical engineer, and a neurobiologist who studies mechanisms of motor control. The authors acknowledge multiple home departments, funding from the National Institutes of Health, and even warn the reader "the authors have applied for a patent on parts of the work described in this paper." Though you will have to take my word for it here, I suggest that these characteristics are entirely typical for research in neural engineering and, more broadly, in much of the work being published in the top "scientific" journals of the present day. In these published works of technoscience, we see crystallized the entanglement of science and society. It is an entanglement, as I describe in chapter 1, that we can understand in terms of existing STS resources: Jasanoff's "bioconstitution" of identities, Galison's "trading zones," Ezrahi's science-democracy linkage, and so on. A philosopher of (techno)science is, thus, confronted with a wealth of objects for normative study, almost all of which elude a solely epistemic focus. As a sort of philosopher's canary, the contemporary scientific literature constitutes a warning that the narrow epistemological approach is increasingly dangerous; some fresh air would do the discipline some good.

Rhesus neuroprosthetics aren't true or false: redeeming normative philosophy of science

The hybridity of technoscience, when we attend to it, can compel action on many fronts. Members of the public may worry about industry's influence on knowledge creation. Scientists may find themselves forced to more explicitly predict the societal implications of their work.[1] The list goes on. But what is the worry relevant to philosophers? In presenting the example from *Nature* and its particular crystallization of co-productive processes underwriting science and society, I am not calling for a complete exodus of scholars from philosophy of science to STS. Such a call would presuppose the very boundaries I am trying to dissolve. I do, however, intend to problematize the inward-looking insularity of philosophy of science in terms that philosophers can under-

[1] Indeed increased demand for accountability to the public may be the only feature of technoscience worth calling novel: see Schiemann (2011)

stand: propositional assertions and published texts. While some clever hermeneutics might allow us to extract just the logical or epistemic content from the 2008 paper, I find it difficult to justify such a narrow focus. Neuroprosthetics, their alleged success or failure, their presupposition of instrumental value to disabled persons, their equal reliance on neuroscientific reasoning and engineering know-how: these are not things we can label "true to nature" or model with probabilistic reasoning. Even if we could extract a hypothesis or theory out of the paper, and judge it from the perspective of philosophy of science, such an extraction would be to ignore philosophy's rich normative potential, as found in ethics, political philosophy, and feminist critical theory. What then, is a philosopher to do with modern technoscientific practice?

The long answer to that question is, essentially, the goal of this dissertation: I call for an account of "good" technoscience rather than just "true" technoscience. In keeping with the traditional normative orientation of philosophy, we need to explain how good science is possible despite its role in the "co-production" (Jasanoff, 2004) of culture and nature.[2] But to do so does has almost nothing to do with *separating* the epistemic from the social. To the contrary, the lesson of co-production is that no epistemological process can be "good" without reference to an accompanying political or ethical trajectory. The implication here is that every epistemological investigation should also be a socio-ethical investigation in some form or another. Such an implication has not gone unaddressed by all philosophers of science. Developments in social epistemology, contextualism, and the "practice turn" suggest an openness to the co-productionist perspective.

By the 1950s, philosophers were already exploring the ways in which the activities of science involved socio-ethical values. C.W. Churchman was one of the first to doubt that scientific investigation is completely described by an account of inductive or statistical confirmation.[3] Scientific inference must be understood, he argues, as a pragmatic search providing means for some prior ends. Shortly after, Richard Rudner's (1953) "The Scientist qua Scientist Makes Value Judgments" took this argument even further. He argues,

> Since no scientific hypothesis is ever completely verified, in accepting a hypothesis the scientist must make the decision that the evidence is *sufficiently* strong or that the probability is *sufficiently* high to warrant the acceptance of the hypothesis. [This

[2] Here, "co-production" refers to the central concept in *States of Knowledge*

[3] See for example Churchman (1948)

> decision] is going to be a function of the *importance*, in the typically ethical sense, of making a mistake in accepting or rejecting the hypothesis.

This feature of hypothesis acceptance, often called "inductive risk", seems to reserve a core role for values in science.

Responses to Rudner and Churchman, even up to the present-day, have called into question the extent to which scientists actually accept their hypothesis when reporting results. Hugh Lacey, for example, is careful to distinguish hypothesis acceptance for the purposes of moving a research project forward and the acceptance of a hypothesis as informative or ready to be applied in non-scientific contexts.[4] These distinctions reflect different pragmatic moments that should not be conflated if we take practice seriously. Moreover, it is possible that "inductive risk" in technoscience requires its own formulation, since technoscience tends to involve claims of control rather than claims of hypothesis confirmation.

Regardless, there's much more to science than hypotheses. Scientific research is often used by scientists and engineers to create or license new technologies and medical treatments; it is often these end products, after all, that are featured most prominently when science appears in media or the popular press. Jasanoff (1990) highlights the political significance of "regulatory science," in which scientific experts negotiate their different disciplinary perspectives through the advisory process. More recently in *Science, Policy and the Value-Free Ideal*, Heather Douglas takes these features of science and, drawing on the arguments of Rudner and Churchman, makes a case that scientists have responsibility to consider the consequences of their errors. She asserts that since scientists are not isolated from society – as the old value-free ideal assumes – they cannot ignore the potential effects of a false positive or false negative result. Contextual values have, and should have, a role in science.

While Douglas and Rudner might oversimplify the phenomenon of hypothesis acceptance, the general point is enough for my purposes. I see inductive risk as only one prominent instance of philosophy trying to accommodate one implication of co-production. There are others worth considering in the scope of this dissertation. Lacey (2012) argues that much of science – specifically science guided by materialist strategies presupposes a particular "valuation of control." This threatens the idea that science is "impartial", and contains no value judgments. Helen Longino

[4] Lacey (2005), see p79

(1990), similarly, has illuminated the inevitable role of background assumptions in enabling scientific evidential reasoning, for better or worse. Based on their understanding of the socio-ethical underpinnings of knowledge creation, each of these philosophers attempts their own reconstruction of the normative philosophical project. I assess a subset of these normative projects over the next few chapters, but for my purposes here, I'd like to foreground in general terms what a philosopher of science should (or should not) do if she takes the processes of co-production seriously.

Despite the efforts of some value-sensitive philosophers of science, the field provides us with many negative lessons. Some common activities therein ignore the co-productionist challenge. Most obviously, a co-productionist scholar is dissatisfied with any attempt to ground scientific truth claims in their logical relation to observation statements. Carnap's famous attempt at this reduction, for instance, can be read as committing such an error. As refracted through present-day philosophy of science, the enduring implication is that the logical relations appropriate to knowledge creation can be spelled out independently of the "external" questions of what the knowledge is meant to do or who gets to decide the relevant objects of study. The lesson of co-production is that knowledge creation does not work that way in practice; epistemic activities (*including* second-order practices of the philosopher of science) are unavoidably entangled with socio-ethical or political content. Careful historians of logical positivism, of course, have already made this point and revealed the political side of the Vienna Circle.[5] But philosophy of science tends to lack this perspective regarding epistemological work of the past and present.

Even recent value-sensitive epistemologies attempt problematic boundary work, trying to preserve some core criteria for knowledge that is untouched by the social. Miriam Solomon's program in *Social Empiricism,* for example, incorporates social, "irrational" dynamics into the production of knowledge – non-cognitive factors can be appropriate when they bring about the right consensus, for instance – but ultimately she assumes that we should assess science's knowledge claims by their anchoring in empirical "vectors." As with Carnap's positivist legacy, there is an unanalyzed foundationalist assumption here that we can talk about the truth without talking about the good, that there is a meaningful standard of knowledge that doesn't involve the socio-ethical or the political. This is an impulse one can find in much of philosophy of science, from formal

[5] See for example, Uebel (2012) "Carnap, philosophy and 'politics in its broadest sense.'"

learning theory and Bayesian epistemologies to causal theories of explanation.

Next steps: the co-productionist approach

There may be indirect or secondary benefits to these sorts of philosophical projects, but to the extent that philosophers want to take co-production seriously, they must realize: good reasoning is not always empirical or logical, but it is always political and always entangled with the socio-ethical. This rule applies to the scientist, to the engineer, and to the philosopher as well. My vision for normative philosophy of technoscience, accordingly, is oriented not towards just truth or objectivity but also towards justice and morality; I call for an account of "good" technoscience, something that we can aim for when we design and assess current technoscientific practices. This reorientation may be disorienting for philosophers of science, but it can be seen as reconnecting epistemology of science to well-honed tools in other areas of philosophy, to ethics and political philosophy. For this reason, I like to refer to this program as a "unifying" or "co-productionist" approach to philosophy of science. In so doing, I see myself as continuing a self-corrective started within philosophy of science long ago.

In order to advance my program for philosophy of technoscience, I will consider previous attempts to walk down the "unifying" path, most notably, that of Helen Longino, Lorraine Code, and Heather Douglas. I begin the project in Chapter 1, where I make the challenge posed by technoscience a bit clearer. There, I draw on science and technology studies to argue that philosophers of science can no longer ignore ethical or political content of technoscience. In Chapter 2, I shift to Helen Longino's "contextual empiricism" and suggest that, while her account could be compatible with a co-productionist outlook, her choice of language is often too ambiguous to fulfill this potential. Hers, however, is not the only promising framework. Heather Douglas and Lorraine Code attempt their own synthesis by subsuming epistemology under the banner of responsibility, which already includes ethical and political responsibility. In Chapter 3, I show how this reorientation also faces difficulties; the scientific practice imagined by Douglas and Code does not adequately capture the actual activities and communities that define present day technoscience.

With these criticisms laid out, I then begin some constructive work. Chapter 4 provides would-be co-productionist philosophers with some relevant social theory and presents my own qualitative research within spaces of neural engineering. I use Chapter 5 to abstract away from these specific qualitative findings and direct attention to new objects, new entities that should be accounted for in any theory of technoscientific accountability; I spell out some desiderata for a theory of "good" technoscience and suggest some tools that can philosophers focus their inquiry. Finally, in Chapter 6, I conclude with a short reflection on the co-productive dimensions of midstream engagement in technoscience, including my own at the CSNE.

References

Churchman, C. W. (1948). Statistics, pragmatics, induction. *Philosophy of Science*, *15*(3), 249–268.

Fetz, E. E. (1969). Operant conditioning of cortical unit activity. *Science*, *163*(3870), 955–958.

Jasanoff, S. (1990). *The Fifth Branch: Science Advisers as Policymakers*. Cambridge MA: Harvard University Press.

——— (2004). The idiom of co-production. In *States of Knowledge: The Co-production of Science and the Social Order* (pp. 1–13). New York: Routledge.

Lacey, H. (2005). *Values and Objectivity in Science: The Current Controversy about Transgenic Crops*. Lanham: Lexington books.

——— (2012). Reflections on science and technoscience. *Scientiae Studia*, *10*(spe), 103–128.

Longino, H. (1990). *Science as Social Knowledge*. Princeton: Princeton University.

Rudner, R. (1953). The scientist qua scientist makes value judgments. *Philosophy of Science*, *20*(1), 1–6.

Schiemann, G. (2011). We are not witnesses to a new scientific revolution. In A. Nordmann (Ed.), *Science Transformed?* Pittsburgh: University of Pittsburgh.

Solomon, M. (2001). *Social Empiricism*. Cambridge MA: MIT Press.

Uebel, T. (2012). Carnap, philosophy and "politics in its broadest sense". In R. Creath (Ed.), *Rudolf Carnap and the Legacy of Logical Empiricism* (pp. 133–148). Dordrecht: Springer.

Velliste, M., Perel, S., Spalding, M. C., Whitford, A. S., & Schwartz, A. B. (2008). Cortical control of a prosthetic arm for self-feeding. *Nature*, *453*(7198), 1098–1101.

1 Confronting Co-Production: Reorienting Philosophy of Science towards the Socio-Ethical

As I stressed in the introduction, 21st century philosophers of science find themselves in a peculiar position. Their discipline carries within it traditional intellectual puzzles dating back to the beginning of the 20th century and likely much farther than that. The ambition of creating a guiding empiricist or probabilistic epistemology of science, for example, persists despite the "turn to practice." We philosophers have available, often a few keystrokes and clicks away, decades of feminist critiques and science and technology studies (STS) that show how deeply science is implicated in the creation and maintenance of power structures – especially state power – in Western democracies. If we were to read them, *Nature* and *Science* now provide us with as many patented proof of concepts as they do confirmed hypotheses or pieces of causal reasoning. One can get a sense of this from the way Fetz' corner of neuroscience grew to include neuroprosthetics. And yet, we philosophers of science continue to ask questions of the form "but how is it knowledge?" and "how can it be true of the world?" and try to delineate the epistemic from the social in ever more subtle ways. In the name of "explanatory power," "robustness," or "causality," philosophers of science relegate politics, values, and power to other disciplines and non-philosophers.

Flipping through a recent issue of *Philosophy of Science* gives a sense of the fields' strangely narrow scope of analysis. There is an article weighing the merit of triangulation; the authors explore principled, abstract reasons for using multiple, independent sets of evidence. Another author seeks to defend Christiaan Huygen's physics in terms of what theoretical work his ideas could accomplish. Another still proposes a new condition to be applied to models of general relativity for the purposes of conceptual clarity. In their own way, each of these authors seems to aspire to help advance Science (or something taken to be science-like) on its own terms. This persistent impulse to get things "right" is a core part of the philosophical project, a strength which sets it apart from many of the other humanities. And yet, these authors externalize many of the same normative dimensions that are externalized by researchers in the natural and physical sciences. They carefully limit their language to the epistemic, despite the fact that no biologist or physicist may ever appreciate that act of self-discipline.

The appropriate intellectual response to this situation, I will argue, is a new orientation for philosophy of science that replaces such rituals of epistemic purification.[1] Rather than imitate the "boundary work"[2] that scientists themselves carry out, performatively cleansing nature of culture, I highlight the need to more directly engage with issues of value within science, with technoscience as an epistemic *and* socio-ethical practice. Only then can philosophers make the most of their commitment to clarifying "the good" and how we ought to achieve it. To make sense of this reorientation, it is necessary to clearly identify the dimensions of science that evade narrowly epistemological philosophy of science. What exactly is left out, hidden, or silenced when we analyze science solely in terms of its reasoning or the logical structures of its textual output? The answer, as presented in this chapter, is made quite evident outside of disciplinary philosophy. STS scholars since the 1970s have developed and honed a whole set of stories and frameworks with which to understand the entanglement of knowledge and culture.

While a comprehensive review of STS resources is not necessary for the philosophical purposes at hand, a few key points will suffice to underpin my prescriptions for philosophy of science. Most significantly, STS scholarship shows that there is nothing wholly unprecedented about the socio-ethical valence of recent neuroscience, as an instance of recent technoscience. Science has always had a significant socio-ethical component that is not easily separable from the creation of knowledge, so they argue. Nevertheless, STS scholarship also suggests that we must pay attention to the *particular*, local links between science and society at different times and different places, especially those closest to us in the present day. These two points come together to capture what philosophy of science has, for much of its history, missed; they emphasize the interdependence of epistemic practices and the social world, highlighting the connections between science and democracy, between funding agencies and lab practice, and so on. By the end of the chapter, I hope to show that philosophers can gain from taking these studies as a starting point, while keeping an eye on "the good." I conclude with some methodological criteria, spelling out this "co-productionist" approach for philosophers of an epistemological *and* socio-ethical persuasion.

[1] I stress here that my proposal is intellectual, but that is not to deny the necessity of corresponding social and institutional components. It would not be effective, for example, to simply start submitting more political work to *Philosophy of Science*. The editors likely would not know how to justify its inclusion.

[2] See Gieryn (1983) "Boundary Work and the Demarcation of Science."

The co-production of science & society and other old news

Is there something new about contemporary technoscientific practice? Does technoscience differ fundamentally from older, "purer" science conducted in past centuries or from the natural philosophy that preceded it? Bruno Latour doesn't think so. He contends that we have never successfully isolated the natural from the cultural. Despite the Western pretense that such a separation is possible – and indeed successful – in the Modern era, "we have never been modern" (Latour, 2012). The world is swarming with socio-technical "hybrids", even if we insist that they exist independently of our contingent aims or our situated perspectives. Take the microbe, for example. It is often assumed to be a fixture of the objective world. Yet, in Latour's analysis, a particular (in this case microbiological) understanding of the natural world only took hold within society because Louis Pasteur made his research relevant to existing social and material realities around him. By carefully "staging" key experiments in agricultural settings, in contexts where the microscopic gained macroscopic and extra-scientific significance, he triggered a restructuring of French society around his lab and his microbes.[3] Successful laboratory work is defined, Latour explains, by its ability to extend its boundaries to encompass the social and political world around it; he calls this power to restructure society "political" in the most basic sense. Laboratory work entails social work.

This feature of technoscience, its translation (successful or not), is philosophically-interesting along at least a couple of dimensions. First, the specific ability for a research program and its proponents to transform communities outside their own provides a new perspective on old philosophical concerns like truth and realism. A carefully told story of technoscientific success, rare as they are, demonstrates the immensity of achievement in the enrollment of people, materials, institutions, and symbols. In this way, obtaining the world's cooperation is not just a matter of logic or correspondence to an external material reality, but rather a feat of equal or greater difficulty. It provides a new lens on epistemic success that couldn't be farther from wishing-makes-true or extreme relativism. Second, we can subject the political changes wrought by successful technoscience to direct evaluation. By making political processes of successful science explicit, we can compare them to our vision(s) of how society should be ordered, whether democratic, totalitarian, or something in between. New publics and citizens, as well as rights and responsibilities,

[3] Described in Latour (1983) "Give Me a Lab and I Will Raise the World"

emerge as we learn about the natural world; consider how the fact of immunity and vaccination is accompanied by new expectations on individual disease-vectors.

Technoscientific investigation of the natural world can also be political in a more direct sense. The process of representing and explaining phenomena can simultaneously dictate a paradigm for who can make legitimate knowledge claims and with what techniques. Shapin & Schaffer (2011) reveal this dynamic in their interpretation of the debate between Boyle and Hobbes. As each natural philosopher battled to establish a standard method for knowledge creation, they also created separate visions for the political structure of Restoration England. Hobbes' *Leviathan* prescribed a causal, "demonstrative" epistemology, analogous to reasoning in geometry. As a solution to the problem of social order, this epistemology left no room for individual judgment; authority comes from above, *a priori*, whether from the principles of the philosopher or the power of a monarch. Boyle's contribution, an experimentalism of leaky air pumps and staged performances, encouraged the observer to witness and decide for himself. The community implied by Boyle's knowledge practices was to be one of a self-disciplined public, lacking any central authority. Shapin and Schaffer, thus, reveal the social and political valence of epistemic practices. They "display scientific method as crystallizing forms of social organization and as a means of regulating social interaction within the scientific community."[4]

Shapin and Schaffer do not see this particular debate as an unusual case of corrupt politically-oriented science. They suggest that in general, "the problem of generating and protecting knowledge is a problem in politics, and, conversely, that the problem of political order always involves solutions to the problem of knowledge." This is a point reinforced and made even more relevant by *Descent of Icarus* (1990), in which Yaron Ezrahi connects political culture of 20th century America with Boyle's reliance on the self-disciplined witness. There are, he explains, deep parallels between the seemingly universal character of science's "visual attestive culture" and the modes of legitimation in liberal democracies, between "stagecraft" and "statecraft."

Just as the experimentalist invites skeptical witnesses, the political actor in a liberal democracy invites the public to judge her actions and claims. In both cases, exposing oneself to an audience serves to depersonalize (politically and morally) actions and claims. And to the extent that the experimentalist or the political actor can portray themselves as constrained by the world of pub-

[4] *Leviathan* p14

licly observable facts, their claims earn legitimation. So it is no surprise, Ezrahi says, that public action in the United States is often allied with scientific and technological projects, which were accompanied by an aura of universality (at least at mid-century). What this amounts to, Ezrahi explains, is not a reduction of liberal democratic governance to the culture of science or vice versa. He is not making a simple causal claim that one led to the other, but rather that the two stand or fall together, as connected practices of scientific and political accountability.[5]

These are only a few of the stories one could tell about the interdependence of epistemic practices (as exemplified in science and natural philosophy) and broader society. I suggest that we can productively understand this type of account in terms of "co-production" as described by Sheila Jasanoff (2004). Her "idiom of co-production" is meant to convey that the natural world is not predetermined, waiting to be discovered and exploited by social actors. And, conversely, the social world is not fixed or logically prior to the natural world. Rather, science, technology, and society underwrite each other's existence, co-producing each other. I read Shapin and Schaffer, Latour, and Ezrahi as revealing particular instances of co-production; they document some of the ways in which the social and the epistemic are interdependent. This relation, as I will suggest later in this chapter, problematizes reductive theoretical positions in philosophy and elsewhere; social determinists and epistemological foundationalists alike have the burden of proof to justify their one-sided analysis of knowledge. But in the meantime, I'd like to consider what co-production can tell us about the characterization of "technoscience" and its place in the scholarly literature.

Recent developments in technoscience

In establishing co-production as a persistent phenomenon, authors like Ezrahi and Jasanoff imply that there is nothing entirely novel about the socio-ethical valence of contemporary science and engineering. In this, they parallel Bernadette Bensaude-Vincent's definition of "technoscience". Bensaude-Vincent (2008) warns that frequently the term does not actually denote a unique neutral description of a contemporary practice; she suggests that it is better thought of as the hallmark of a dialectical move away from traditional conceptions of science. "Technoscience", she suggests, is used as a polemical tool for which "the only stable connotation is the break with

[5] see especially *Descent*, p75

modernity or at least modern mythologies." All science is technoscience, these authors might say, and that is an assertion that controverts the traditional ideal of pure and disinterested reasoning on the given. It is important, however, to not overemphasize this point – a denial of novelty – and, as a result, miss more subtle developments in technoscience that are relevant to our present situation.

There has been a growing sentiment in the scholarly literature that there is something troubling (or at least worth studying) in recent history of science and engineering. STS scholars and even some philosophers use the word "technoscience" to emphasize changes in science and engineering since the second world war. But rather than ignore this scholarship as "old news" or simply "co-production by another name," I would like to highlight the way this literature gives a more detailed understanding of co-production in our own era and gestures at its progression from older socio-technical arrangements. In doing so, I will focus on roughly three aspects of "technoscience": its implicit epistemology and ontology, broader changes in institutions, and overarching societal dynamics.

At the micro-analytical scale, scholars have looked to published pieces of recent technoscience to reverse engineer its characteristics. Bensaude-Vincent et al. (2011) take an ontological approach to analysis. There they reveal objects of technoscientific study to be "performative" and "value-laden"; technoscientific research "engages the things, participates in their agency and establishes their value" in terms of some functional property like reflectivity or conductivity. And unlike the traditional ideal of science, there is no careful attempt to separate artifacts of human intervention from the "pure" features of nature. Bensaude-Vincent and her co-authors recognize, however, that this ontology, as an ideal, might operate simultaneously within the same field of research as more "traditional" scientific ideals, depending on the context. In other words, a technoscientific actor might alternate between representing features of the natural world and manipulating the functional properties of human-made objects.

Alfred Nordmann (2012), similarly, has his own epistemological take on these features. He observes that technoscientific publications are often oriented around claims of control, rather than around claims of evidence. Technoscientific knowledge is characterized by the proclamation of a new capability, which is made compatible with existing "epistemic knowledge" only indirectly. But despite his observations, Nordmann is also careful not to make any claim that these features is entirely unique to present-day research or that they are necessary and sufficient conditions for research program to be technoscience. Thus we see, if we read publications from nanotechnology, synthetic biology, and similar areas of converging technologies, that much of recent science and engineering has explicitly embraced an ontology and epistemology that is not limited by the idea of science as testing hypotheses about the natural world.

Zoom out to the level of institutions, and one finds more indications of change. Ezrahi (1990) detects a set of recent changes in American political culture and in science. He notes that, before the 1960s, the National Science Foundation (NSF) could fund research more or less independently of congressional approval. But following the requirement for annual congressional reauthorization of federal spending, the NSF leadership – like those in any other federal program – found themselves forced to justify their expenditures in terms of short term benefits to society. Ezrahi asserts that the authority of internal professional scientific norms was weakened as the NSF responded: "[finding] it necessary to protect its base of political-public support by redefining "basic research" broadly enough to include also categories of applied research and, in particular, research directed at helping to solve practical economic and societal problems."[6] American scientists were, for the first time, exposed to public appraisal and were drafted to work towards their collective good. Since that shift, Ezrahi asserts, science in America is used as "a resource in competitive ventures, such as economic and industrial growth or military conflict, and less as an intrinsically valuable universalistic cultural activity."[7] It is this instrumental attitude that is institutionalized by technology transfer offices at public universities, the Bayh-Dole Act of 1980, and in entrepreneurial seminars (e.g. Neuroventures 101) offered to young scientists.

Peter Galison, more recently, echoes Ezrahi's observations about the instrumentalization of post-WWII science. In "Trading with the Enemy," Galison (2010) describes how scientific practice is altered by its recruitment into application. He tells a story of how forces external to science - war - drove high theory out of isolation and into zones of technological development. During the height of World War II, he explains, theoretical physicists on both sides of the conflict found themselves thrust into an unlikely partnership with radio engineers. By the end of the war, elements of circuit design started to appear in purely theoretical work in physics. The techniques used by radio engineers to simplify circuit design, a low-prestige activity, became central to the formulation of

[6] *Descent,* p256
[7] ibid. p278

renormalization theory. Somehow, the "rules of manipulation" used by radio engineers had found a way into seemingly pure theoretical pursuits. For Galison, this historical episode shows how social forces can push disparate communities together, even communities of drastically differing power and prestige, such that members begin to talk to one another and trade ideas developing a hybrid "interlanguage." This effect of trading zones is remarkable, not because it is rare, but rather because it is common. Galison hints that we can find it throughout 21st century society, in the interactions of science and government, in the shared labs of nanotechnology, and in scientific startup culture.

Finally, at the highest level of analysis, scholars have stressed new and significant interplay between the development of technoscience and changes in global social organization. Jasanoff (2003) understands these broad changes as "constitutional," asserting that the "basic ordering commitments" of a society are not limited to legal documents or the US Constitution. To the contrary, genetics, informatics, nanotechnology, climate science and other technoscientific practices must be recognized as active forces in "the redefinition of self, identity, and community; the appearance of the consumer as a political agent, asserting rights claims against commerce and industry; and the certification of 'global' knowledge for use in supranational governance."[8] This perspective on recent technoscience – in tension with Latour's simple assertion that "we have never been modern" – suggest that co-production is not a foregone conclusion or a mundane fact, but an active and consequential process, shaping our global future. It is this insight that is emphasized by Gibbons et al. (1994) when they herald the dawn of a new, Mode 2 form of knowledge production, when they draw our attention to unique features of the "coevolution of science and society" in the last few decades.

Both Jasanoff's constitutional lens on technoscience and the new "modes" identified by Gibbons and Nowotny et al. compel scholars to move beyond *historical* case studies of co-production and pay attention to the simultaneous emergence of knowledge, artifact, and social order *in the 21st century*. The entails that we investigate visions of the good life implicit in brain-computer interfaces and chimeric organisms. It means grappling with the ways that the idea of the person is re-negotiated by IRBs and bioethics committees as they work to define the limits of the human (and non-human) and the scope of responsibility. It means we must acknowledge the fact that

[8] "In a Constitutional Moment," p157

scientists and engineers often work within the same labs, that patenting often takes priority over publishing, and that now (more than perhaps ever before) consumers, sociologists, and members of the public are calling for an "opening up" of the "protected spaces" of technoscience, demanding representation even at the benchside.[9] Only when we expand the scope of our analyses can we, as scholars, assert that we have learned the lesson of co-production, as highlighted by Shapin and Schaffer, Ezrahi, and others.

Good or just true: technoscience's challenge to philosophy

It is my contention here that the co-production at work in contemporary technoscience presents a broadly unmet challenge and an invitation for philosophers of science. STS scholarship identifies a whole series of new relationships and entanglements that require critical normative analysis, yet science and technology studies scholars have responded unevenly to this challenge. Sometimes a moment of co-production is simply presented for the reader to consider. Shapin and Schaffer's *Leviathan*, for instance, is given as a state of affairs, but we are never asked if we should aspire to create a social order based on Hobbes or on Boyle. Other STS scholars are apt to push a bit farther and problematize particular moments of co-production. Jasanoff for example, in "Judgment Under Siege," highlights and questions the features that make our current knowledge-society system cohere. Still, she does not reflect on what would make these assemblages worthy of our praise; how can we pair our re-examination of expert scientific judgment with our own judgments of "the good"? Philosophers may be well-positioned to answer this question, if they re-orient their inquiry towards dynamics of knowledge/power.

This proposal, that philosophers should take co-production as a starting point, will strike some as obviously incorrect; many of us possess a reflex-like response to any call for sociological or political understanding of science. Faced with extreme social constructivism, whether embodied by Edinburgh School sociology of scientific knowledge (SSK) or by climate change deniers in a deconstructive mode, philosophers of science can stomp their feet and assert that it is the world that constrains our best scientific knowledge, "bad" or "contaminated" science aside. Pinch and Bjiker's "empirical programme," for example, defines scientific truth in terms of "closure," a result

[9] To borrow language from Rip (2011)

of inherently social process of ruling out alternative interpretations of the world. To many philosophers, such a reduction appears toxic; if truth is nothing more than the contingent outcome of social processes, how are we to understand our own ongoing practices of inquiry? Wouldn't scientists find it self-contradictory, in some sense, to understand their personal research as simply input into an arational system of entirely social forces? Philosophers of science can use these sorts of worries to rationalize their own narrowly epistemological focus. As long as technoscience is presented in reductive terms, it cannot pose a clear challenge to philosophy of science.

The philosophical reaction to sociology of scientific knowledge (SSK), social construction of technology (SCOT), and other seemingly reductive frameworks deserves its own chapter, arguably, but here they serve as foils to the co-productionist approach, which does not allow philosophers their well-honed reflex. The idiom of co-production is significant for its avoidance of mere deconstruction, reduction, and linear explanations. This feature is not accidental; Jasanoff explains that the label "co-production" is an explicit attempt to counteract "thin" readings of science and technology studies, and to remedy a long-standing dissatisfaction with the misleading connotations of "social construction."[10] The latter, for instance, might suggest that scientific representation is just an elaborate social game, with no meaning apart from its use by actors to advance their social interests. Jasanoff's idiom, on the other hand, reminds us that knowledge—including scientific representation—is foundational for *any* human action and social order. The idiom thus suggests that knowledge practices and their products have meaning that is broader than their strategic use by opportunistic individuals.

For philosophers, this means that a robust normativity about knowledge is not precluded by a sensitivity to co-production. Quite the opposite is true, in fact. But that does not help us avoid the "challenge" to be answered; there is a sense in which normative epistemology must be reimagined within a more encompassing framework. Jasanoff suggests the direction that this might take: "Co-productionist accounts take on the normative concerns of political theory and moral philosophy by revealing unsuspected dimensions of ethics, values, lawfulness and power within the epistemic, material and social formations that constitute science and technology."[11] Just as I hinted in the introduction, co-productionist philosophy of science will have to accommodate the moral and political elements in technoscience. How this might happen, and whether previous

[10] *States of Knowledge*, p43
[11] *States of Knowledge*, p4

value-sensitive epistemologies are adequate, will be a core question as I proceed in this dissertation.

Co-productionist method

We should strive to develop tools within philosophy of science that both provide normative leverage and acknowledge the fact of 21st century technoscience. In other words, we must engage its entanglement with non-epistemic elements. Taking the perspective of co-production vindicates such a project. But what would such an approach look like? In the next few chapters, I will consider some promising frameworks, specifically those of Lorraine Code and Heather Douglas, as well as Helen Longino. But here, I would like to lay out some preliminary desiderata of a co-productionist philosophy of science, some virtues that should define future work. Returning to the example of Rhesus monkeys, as presented in the introduction, will help me to make three methodological points.

Recall that the interesting contrast between the paper in 1969 and 2008 was the emergence of new types of bold claims, both ethical and political. In response to this, philosophers of technoscience should develop a sensitivity to all categories of "validity claims";[12] just as modern technoscientific actors are well-practiced in making political, moral, and epistemic claims, so should philosophers be able to find validity claims of all sorts in published work and other official discourse. As the 2008 paper demonstrates, it takes very little work to find claims that are suitable for (bio)ethical, political, and of course epistemic, analysis. We could ask, respectively, if the normalizing function of a neuroprosthetic is beneficent, if society should require the individual to cure her own disability, or if Rhesus monkeys are a suitable model organism for understanding plasticity and "embodied" motor control in the human brain. In practice, these questions are entangled.[13] And their entanglement justifies, I hope, my call for a co-productionist approach to philosophy of science. This even handed sensitivity to claims beyond the epistemic might be achieved by value-sensitive epistemologies, like that of Helen Longino (the subject of Chapter 2), as well as Heather Douglas' responsibilist framework (the subject of Chapter 3).

[12] to use Habermas' terminology

[13] Consider that we could not define "suitable animal model" without reference to some pragmatic aim of the associated research. And even if the researchers did settle on some "internal" definition of suitability—maybe a genetic metric of similarity—their decision has implicit political assumptions about the role of science in society, performing a particular arrangement between neural engineering and its publics.

The second methodological criterion takes Shapin and Schaffer's argument seriously; philosophers of technoscience must keep the political implications of epistemology firmly in mind. This is a key lesson of co-productionist scholarship. Whenever an epistemic claim is made within technoscientific discourse, the philosopher should keep in mind that it likely represents some restructuring of society (a la Pasteur) or a profoundly political paradigm of legitimation (as in Boyle's laboratory). The former is more obviously at work within neuroprosthetic research; the authors cannot help but project future social effects of their findings. Yet one can, nevertheless, ask why videos of monkeys feeding themselves is to be taken as credible evidence and question who that standard of evidence effectively disenfranchises or silences in neuroprosthetic development.

Third, and finally, co-production applies reflexively to the philosopher herself; philosophers of science cannot impose an epistemology or a definition of knowledge onto a practice without recognizing the political implications. To impose, say, an empiricist or a feminist-critical criterion on a scientific practice is not a return to neutrality but is itself a political act. Conversely, every ethical or political framework has epistemic implications, for who can make knowledge claims and how. Though philosophers rarely have the power to impose their normative frameworks on non-philosophers, any power that we do accrue must be used responsibly. That is to say, if philosophers fulfill their desire to become socially engaged and scientifically-respected, then our own discourse becomes part of the co-production of social order and the natural world.

In phrasing these criteria in cognitive-theoretical or intellectual terms, I am indeed asking philosophers to "think different."[14] But that proposal does not reduce to just thinking differently. There are non-trivial social and institutional circumstances that have led philosophers to write and speak in the way that they do. Similarly, the creation of specializations within our field did not occur by theoretical fiat. In the face of ongoing disciplining dynamics, the co-productionist philosopher will have to create new knowledges simultaneously with new forms of life for philosophers. We might need new places to publish, which may not be academic journals. We might need new scholarly alliances, which could take us out of the department. Exposure to technoscience and collaboration with STS scholars was essential for my own work, yet it happened more or less by chance. I did not plan to be one of the very few graduate students who are allowed (and able) to

[14] A phrase canonized by Steve Jobs' Apple

spend a significant portion of their PhD work adjacent to technoscience. If my experience is any indication, then doing co-productionist philosophy will be equal parts intellectual and social.

References

Bensaude-Vincent, B. (2008). Technoscience and convergence: a transmutation of values? *Summerschool on Ethics of Converging Technologies, Dormotel Vogelsberg, Omrod/Alsfeld, Germany*.

Bensaude-Vincent, B., Loeve, S., Nordmann, A., & Schwarz, A. (2011). Matters of interest: the objects of research in science and technoscience. *Journal for General Philosophy of Science*, *42*(2), 365–383.

Ezrahi, Y. (1990). *The Descent of Icarus: Science and the Transformation of Contemporary Democracy*. Cambridge MA: Harvard University Press.

Galison, P. (2010). Trading with the enemy. In M. E. Gorman (Ed.), *Trading zones and Interactional expertise: Creating New Kinds of Collaboration*. Cambridge MA: MIT Press.

Gibbons, M., Limoges, C., Nowotny, H., Schwartzman, S., Scott, P., & Trow, M. (1994). *The New Production of Knowledge: The Dynamics of Science and Research in Contemporary Societies*. London: Sage.

Gieryn, T. F. (1983). Boundary-work and the demarcation of science from non-science: Strains and interests in professional ideologies of scientists. *American Sociological Review*, *48*(6), 781–795.

Jasanoff, S. (2003). In a constitutional moment: science and social order at the millennium. In *Social studies of science and technology: Looking back, ahead* (pp. 155–180). Dordrecht: Springer.

——— (2004). The idiom of co-production. In *States of Knowledge: The Co-production of Science and the Social Order* (pp. 1–13). New York: Routledge.

——— (2006). Judgment under siege. In S. Maasen & P. Weingart (Eds.), *Democratization of Expertise? Exploring Novel Forms of Scientific Advice in Political Decision-Making*. Dordrecht: Springer.

Latour, B. (1983). Give me a laboratory and I will raise the world. In Knorr & Mulkay (Eds.), *Science Observed: Perspectives on the Social Study of Science*. London: Sage.

——— (2012). *We Have Never Been Modern*. Cambridge MA: Harvard University Press.

Nordmann, A. (2012). Object lessons: towards an epistemology of technoscience. *Scientiae Studia*, *10*(spe), 11–31.

Rip, A. (2011). Protected spaces of science: their emergence and further evolution in a changing world. In A. Nordmann & M. Carrier (Eds.), *Science in the Context of Application* (pp. 197–220). Dordrecht: Springer.

Shapin, S. & Schaffer, S. (2011). *Leviathan and the Air-Pump: Hobbes, Boyle, and the Experimental Life*. Princeton: Princeton University Press.

Chapter 2 - Contextual Empiricism, Liberal Utopia, and the Question of Philosophy's Expertise

What are philosophers doing when they prescribe a particular epistemology or procedure for science? The simple answer might be that they are describing how a knowledge practice should work, which may suggest ways to change existing practices or to evaluate suspect claims to knowledge. I once heard a philosopher of science say that they were "just trying to devise better epistemologies for science." By now, the reader should know that I would not allow that statement to pass for an answer. The idiom of co-production suggests that there is also a political and ethical side to every epistemology, especially if they are part of human activities outside of the pages of philosophy journals. Why? Almost immediately an epistemology excludes certain people and ways of speaking from the processes of knowledge production. These excluded individuals are not usually the insane or demon-possessed but are intelligent people who do not appreciate their diminished standing. Even in the news today, I saw an open letter to the *British Medical Journal* from researchers who disagree with the editorial decision to preference quantitative methods over qualitative. And it doesn't take much work to think of even more consequential examples, like the fight for epistemic credibility between AIDS treatment activists and the 1980s biomedical establishment.[1]

The connection between epistemology and politics has not been lost on all philosophers of science. Philip Kitcher, Helen Longino, Karl Popper, and even the Vienna Circle (according to some accounts) explore some of the ways in which our epistemic ideals have implications for social order and vice versa. Longino, in particular, will be the focus of this chapter. I will examine how she grounds her prescriptions for science in a sort of liberalism and whether her account is compatible with co-productionist scholarship. But there is also a pressing question beyond the co-productionist worry: with what authority can a philosopher propose a set of epistemic-political ideals? Just because Longino or I have a vision of a well-ordered society doesn't mean that it deserves any uptake among our fellow citizens; any claim to a philosophical brand of expertise needs to be elaborated. I will pursue these issues in turn. I argue that i) she pursues a co-productionist strategy, but problematically preferences epistemological registers and ii) that her authority as a philosopher rests on an undeveloped idea of representation.

[1] As described by Epstein (1996)

Hobbes and Boyle return

What might Helen Longino say about the co-productionist idiom? Part of an answer can be found in *The Fate of Knowledge* (2002). There, she takes a few pages to directly consider the argument posed by Shapin and Schaffer in *Leviathan and the Air Pump*, which happens to be one of the key examples of co-production. Longino admits that the Hobbes-Boyle story does an excellent job illustrating the "parallels between conceptions of political authority and cognitive authority," but she is not convinced by their "causal argument."[2] If you've read the book, you may be confused by this statement, because Shapin and Schaffer do not make any explicit reference to the causes of belief. The closest they get, perhaps, is when they say that they are interested in "how and why certain beliefs were taken as true."[3] Philosophers will not be satisfied with such ambiguous wording, but I have a hunch that Longino's interpretation may be too extreme. I will explain my side of the story soon enough, but first it would help to understand the worry in *Fate of Knowledge.*

Longino's phrasing ("causal argument") makes sense if we notice that she interprets *Leviathan and the Air Pump* as a canonical text of Strong Programme Sociology of Scientific Knowledge (SSK), as illustrated by Barnes and Bloor's (1982) "Relativism, Rationality, and the Sociology of Knowledge." Unfortunately, the ground is treacherous here. SSK is characterized by so many different formulations and theses, examples and applications. Longino herself provides a strong and moderate reading of Barnes and Bloor (1982). She suggests, for instance, that their stance on causal explanation might be: "for any given belief, regardless of whether it is true of false, rational or irrational, its being taken to be true or reasonable in a context requires a causal explanation." That's the moderate reading, but maybe Barnes and Bloor intend something more: "for any given belief, regardless of whether it is true of false, rational or irrational, its being taken to be true or reasonable in a context requires a causal explanation *that makes no reference to the truth of or reasons for the belief.*" Shapin and Schaffer may be read as committing an error somewhere along this dimension, making an illicit causal explanatory connection between Boyle's beliefs and something social.

Longino concludes that, while Shapin and Schaffer provide a nice story about the parallels between cognitive and political authority, they fail to draw any causal arrow, to show that Boyle's

[2] *Fate,* p20
[3] ibid. p14

success in establishing a program of experimentalism is explained by a social interest in a "parliamentary form of life." Why not? Part of the problem, according to Longino, is that Shapin and Schaffer haven't established the appropriate counterfactual (as might be expected for causal arguments). They should have but did not adequately show how Hobbes' preferred political system could have led to a viable non-experimental epistemology. We don't have much to go on, of course, since Hobbes' methods where never widely adopted or institutionalized like those of Boyle. It is important to note, here, that our possible dissatisfaction with their causal argument should not be grounded in an unreflective foundationalism; it would be begging the question to rule out the viability of Hobbes' program simply because we sympathize with anti-authoritarian experimentalism. Longino is not making this mistake.[4]

Besides, she has a more fundamental criticism to make. Longino seems to reject the significance of the explanation that (in her words) "the political dimension of Boyle's experimental method provided an incentive to adopt it."[5] Even if we allow Shapin and Schaffer's example, it does not, she says, "lead to radical relativism or render cognitive explanations moot." Notice that her language of "incentive" treats political commitments as something that impinge on knowledge like a coercive force or a cynical strategy, interfering with or replacing rationality, all "how" and no "why."[6] This gloss of Shapin and Schaffer seems mistaken. At worst, it oversimplifies how knowledge practices and politics interact. Deciding who can credibly interpret experimental results can be a reasonable, even rational process, but the politics remain. At best, Longino's gloss misreads Shapin and Schaffer as social determinists, as the strongest of the Strong Programme.

Latour's (1990) review of *Leviathan* echoes Longino's worry but takes a slightly more charitable angle. He asserts that the authors tend to "treat society as more transcendental than nature," (i.e. Longino's complaint) but not due to a prejudice in favor of social-causal explanation. Latour points out that the objects of Boyle's experiments, Nature's contributions, weren't at issue between Hobbes and Boyle; the controversy was over "the management of experiment," a nominally socio-political issue. The role of the natural world, actants, and their agency is left unresolved.[7] While Latour would presumably have preferred a full actor-network account of the his-

[4]At least not in the passage described here. Her own program of "contextual empiricism" would rule out Hobbes' program, so the reader has to carefully sort out which of her claims depend on her preferred ordering of science.
[5]*Fate*, p37
[6]ibid. p37
[7]It is no coincidence that this issue happens to be at the center of most of Latour's work.

torical episode, he still praises the book for its theoretical lesson, namely, that science studies can reveal the artificially bounded arenas of politics and epistemology as "two sides of the same coin," as a "dual" and simultaneous "invention" at the dawn of modern scientific society.[8] As in Jasanoff's formulation of co-production, we can understand this without giving priority to things-in-themselves and Nature, on the one side, or social order and politics on the other. More crucially, I want to stress that the conclusion drawn by Latour is not a causal one. The issue at hand is not what explains what. A co-productionist reading of *Leviathan* reveals that the separation between natural order and social order is an achievement made by actors, and in this case, by Boyle the "scientist."

It remains *entirely possible* to accept this conclusion and treat individual scientists as reasoning, rational beings with some degree of agency (as opposed to points being pushed around by social vectors). There is no need to deny the explanatory force of cognitive accounts that appeal to what's in actor's heads. I hope this is exhibited in my qualitative work in this very dissertation; we can juxtapose epistemology with normative politics not with anormative "social." What is impossible, in the co-productionist mode, is to treat questions of social order apart from questions of knowledge. It's a shame that Longino did not focus on this aspect of *Leviathan*, since it arguably has great significance for her prescription for science, "contextual empiricism," and echoes her own "non-dichotomizers way" in *Fate of Knowledge*. Accordingly, the question moving forward and in the remainder of this chapter is whether Longino has settled on a co-productionist solution despite her rejection of Shapin and Schaffer. Does she consider "the social" as a normative question rather than an anormative force of nature? Does she ground her epistemological recommendations in a specific political vision for society? In the next section, I suggest that her arguments needs some re-interpretation in order to be read in a co-productionist way.

Longino's liberal epistemology

In keeping with a co-productionist mode of philosophy, Longino combines her value-sensitive epistemology ("critical contextual empiricism") with a political commitment to a Millian liberalism. She thus can be read as proposing a particular co-productionist solution. Unfortunately, she still engages in some problematic boundary work, separating issues of knowledge from is-

[8] Latour (1990), p11

sues of ethics and politics. As I illustrate some of these problematic moments, the reader should not jump to conclusions about their origins. Longino likely had substantive reasons to write her two monographs as she did, facing a very conservative audience in philosophy of science. So moving forward, there is no need to theorize some intellectual failure or personal neglect on her part. Philosophy, like science and any other practice, confronts its practitioners with preexisting structures and entrenched cognitive assumptions.

Bringing philosophy of science into feminism[9]

In *Science as Social Knowledge* (1990), Longino starts with some bold philosophical goals, motivated by puzzles from philosophy of science. She proposes an account of scientific reasoning that takes the middle way between positivism and holism. With it, she hopes to avoid both underdetermination and unbounded relativism. To understand her account, it is useful to spell out the weaknesses of each alternative.

Logical positivism, she worries, cannot provide a satisfying picture of confirmation. She notes that most scientific theories are not mere generalizations of empirical statements; theories often contain non-observational terms that cannot be straightforwardly connected to simple empirical statements. This logical gap leads to instances of underdetermination. To use her example: perceptions of red bumps on someone's torso might support a hypothesis of measles just as well as it does a hypothesis of health, based solely on the testimony of one's senses. Historical cases of underdetermination illustrate the same point, as when the same set of data can be taken to support both Ptolemaic and Copernican theories of motion. Longino suggests, on the basis of this inferential inertia of sense data, that the positivist vision of science cannot be the whole story; science cannot consist solely of the logical construction of theories from observation statements.

Holism, on the other hand, has the resources to explain these cases of underdetermination. Evidence is connected to potential hypotheses by the totality of our other beliefs. It is laden by our other theories and beliefs about measles and the meaning of red bumps in biological contexts, or by our Aristotelian cosmology, and so on. As Kuhn would assert, we can only see the world through our theories. This thesis of theory ladenness is useful, Longino observes, because it al-

[9] This subtitle was inspired by Steve Fuller's (1993) observation that Longino (1990) spends much more time explaining traditional philosophy of science than she does bringing feminist thought into philosophy of science. As I hinted above, this choice was likely shaped by the culture of philosophy of science.

lows holists to explain why one state of affairs or one set of evidence can support two radically conflicting theories. But simultaneously, holism implies that we can never specify evidence or its meaning independently from the hypothesis that it is meant to test. Because of this vicious circularity, scientists can only jump irrationally from one theory to another; comparisons of empirical support are just not possible.

In order to avoid these two extremes, Longino asserts that we must maintain a bit of both. She spells out this middle way, following Mary Hesse in *Revolutions and Reconstructions in the Philosophy of Science*, by pointing out a more limited form of theory ladenness. It is logically possible, she observes, that the evidence brought to bear on a particular hypothesis is theory laden *but not by the hypothesis in question.* In the measles example, I connect the perception of red bumps to a diagnosis of measles via a set of background assumptions. In order to test that I have measles, I must rely on the assumptions about the definition and appearance of measles, assumptions that my perceptions are not mere hallucinations, and so on. These background beliefs, Longino asserts, are "enabling conditions" for scientific inference, beliefs "in light of which one takes some *x* to be evidence for some *h*.[10] But I need not rely on the hypothesis that I am testing, i.e. that I do indeed have measles.

Based on this possibility, Longino proposes that evidential reasoning is a three term interaction between evidence, hypothesis, and background assumptions. Background assumptions function to explain why it is that evidence often can support conflicting hypotheses, but without asserting that our beliefs form one giant conjunction. Thus, Longino's account of reasoning maintains both the positivist intuition that we can specify evidence independently from the hypothesis being tested *and* the holist insight that theory ladenness (in some form) is inevitable. It is value-sensitive because it recognizes that many our background assumptions are a function of our cultural values, social position, or worse, of our idiosyncratic "biases."[11]

Longino's account of evidential reasoning is primarily descriptive, but it leaves room for a complementary normative framework for scientific knowledge. Longino's attempts to fulfill this potential by providing a social account of objectivity. She tries to explain how it is that science can (and does) correct for subjective idiosyncrasies of its participants. She does so in terms of background

[10] *Science as Social*, p44.

[11] In general, I try to avoid the word "bias," because it implies that there is some non-biased perspective that we are trying to locate. For Longino, the "unbiased" perspective is the intersubjective perspective, but that preference of course requires its own arguments and ethico-political commitments.

assumptions:

> As long as background beliefs can be articulated and subjected to criticism from the scientific community, they can be defended, modified, or abandoned. As long as this kind of response is possible, the incorporation of hypotheses into the canon of scientific knowledge can be independent of any individual's subjective preferences.[12]

This passage is telling; it illustrates that Longino thinks of objectivity as *i*) a species of intersubjectivity and *ii*) an absence of individual, idiosyncratic assumptions in reasoning. The goal? A maximally intersubjective canon of scientific knowledge.

Moreover, since objectivity is tied to the norms that are active in a community, it comes in degrees. Depending on the extent to which criticism is integrated into a practice, the community is more or less objective. Longino spells out four norms that must be active in order for criticism to serve its objectivizing role:

> 1) there must be recognized avenues for criticism of evidence, of methods, and of assumptions and reasoning; 2) there must exist shared standards that critics can invoke; 3) the community as a whole must be responsive to such criticism; 4) intellectual authority must be shared equally among qualified practitioners.[13]

Only when these four conditions are met can we say that a community includes "transformative criticism." Longino's vision of a well-ordered scientific enterprise, thus, starts with her account of reasoning and moves to prescriptions about how to maximize objectivity.

Mill and the marketplace of ideas

For the unsuspicious or charitable mind, the program spelled out in *Science as Social Knowledge* seems eminently reasonable. Who could possibly object to the removal of bias from our knowledge practices? Sure, there is a lot of elaboration needed for the four prescribed norms, but for many readers in Western democracies (and perhaps elsewhere) these suggestions appeal to a deep desire for equality among citizens and for substantive processes of collective deliberation and critique. We are at our best, it seems, when we genuinely listen to one another and let argument settle our beliefs rather than appeals to authority or to personal dogma. Science, of all practices, should embody these lofty ideals. The fact that Longino's solution solves the old philo-

[12] ibid., see p74
[13] ibid., see p76

sophical chestnut of underdetermination only adds to this attractive vision for the ordering of science. I caution that this proposal, worthy our assent or not, should be seen for what it is, both epistemic and ethico-political.

The norms proposed by Longino have a long history in political theorizing, especially of the liberal variety. Most directly relevant is John Stuart Mill's call for a "marketplace of ideas." Mill himself does not use the metaphor—it appeared in political discourse sometime after his key writings—but he gives the idea its fullest treatment in *On Liberty* (1859). There, it is immediately clear that Mill is possessed by a single fear, the illegitimate control of persons by concentrations of authority, whether in the government or in the majority. Coercion of the individual, he worries, is increasingly common, acting on the body and the mind: "There is also in the world at large an increasing inclination to stretch unduly the powers of society over the individual, both by force of opinion and even by that of legislation."[14] He stresses that it is not enough that many theorists are already committed to the idea of liberty; he challenges us with a philosophical problem of practical import: "how to make the fitting adjustment between individual independence and social control." And this "adjustment" requires us to tackle more than the obvious cases of physical control or threats of violence. What would it mean for a person to have freedom of opinion and thought, "on all subjects, practical, or speculative, scientific, moral, or theological"?

In the beginning of the first section, "Of the Liberty of Thought and Discussion," he poses a hypothetical. What if society was unified in opinion in opposition to a single person? Would they be justified in silencing or coercing that individual? He warns that to do so hurts society as much as the individual:

> If the opinion is right, they are deprived of the opportunity of exchanging error for truth: if wrong, they lose, what is almost as great a benefit, the clearer perception and livelier impression of truth produced by its collision with error. [...] We can never be sure that the opinion we are endeavouring to stifle is a false opinion.

In the face of limited human certainty, Mill is prescribing a social epistemology that pits ideas against one another. As in Longino's "contextual empiricism," minority opinions and their proponents give us a chance to test our beliefs, to subject them to critique from diverse perspectives. At his most extreme, Mill stresses that truth and certainty can only be accessed in this way. If we face opponents, we must listen to them, and if there are none to face, we must imagine them.

[14] p29

After hearing these arguments, it may be tempting for the reader to ignore Mill's overarching commitment to liberty and focus myopically on truth, but Mill's aversion to doctrinal authority or enforced opinion is definitive for his epistemology. There is no god's eye view from which Mill or anyone can define what arrangement of knowledge practices is most truth conducive. Accordingly, Mill has to ground his idea of truth in his own primary political ideal: liberty, which is itself valued in utilitarian terms. He says, for instance, that we must not assume the mantle of "the judges of certainty" and that we must allow the "fair play" of ideas. The prescription is as much about politics as it is about "truth." Accordingly, it would be misleading for a Millian epistemologist to attack the authoritarian-dogmatist by merely asserting that their hierarchical way of life is not truth-conducive; instead, the Millian should say that an authoritarian definition of truth does not allow desirable forms of life, namely a society that is flexible in the face of change and allows individuals their own cognitive agency.

The form of life we find implicit in *On Liberty* has had lasting popularity, partially due to lasting fears of illegitimate authority. In the mid-20th century, liberal ideals of free human action and thought were tied directly to science itself. Visvanathan (1997) argues that as enthusiasm for the free market began to wane, science and its orientation towards knowledge became the new model for liberalism:

> In the discourses of university dons, science was the model of *communitas.* The Republic of Science was deemed an open society, sustaining a creative tension between individual initiative and collective truth. In this more liberal world, the scientific method was substituted for the invisible hand and Popper and Polanyi became the Adam Smiths of this new regime.[15]

The invisible hand thus moved from economic problems of distribution to the realm of proper belief. In the aftermath of global war and the Great Depression, we also see Merton (1942) first proposing his analysis of the normative structure of science, originally titled "A Note on Science and Democracy." There, he stresses how science allies itself with the central values of liberalism, including equality of intellectual authority and open trade of ideas. These connections should be expected, if we are convinced by the argument in Ezrahi (1990). He suggests that it was Boyle's experimentalism that provided democratic governments with the tools and methods they needed to legitimize their action in the eyes of the public.

[15] Visvanathan (1997), p146

I take this brief moment to situate liberal science in its historical context for a specific purpose. I want the reader to notice that the preference for liberal norms responds to our desires and fears for society. In the mid-20th century, the threat of totalitarianism, of Nazi science, and global crisis provide substantive reasons to argue for an ordering of science that mirrors and supports liberal-democratic governance. The point here *is not* that Popper's *The Open Society and Its Enemies* (1950), for example, is an outdated product of its time. Neither do I intend to suggest that "the social" somehow coerced him and other thinkers to propose anti-totalitarian or non-communitarian models of knowledge production. Such a reading would undermine liberal theorists as agents or as persons that respond to the same hopes, fears, and reasons that fill our own everyday experience, including philosophical writing. Regardless, the primary sources show no such thing.

Popper, for one, seems quite aware of the connections between politics and ways of knowing when he criticizes the totalitarian *Republic,* and suggests we read Plato as a sincere advocate of justice held back by his restrictive epistemology. The theory of Forms led Plato to "build up a political science" and "opens a way, in the social realm, towards some kind of social engineering; and it makes possible the forging of instruments for arresting social change."[16] Popper seems horrified by that outcome and invites us to pursue a different pairing. Rather than seeking the Forms, we should ground our beliefs in critical testing among free and equal individuals. And rather than assembling an immutable totalitarian edifice, we should modify society in a "piecemeal" fashion, allowing individuals to experiment with different forms of life. For him, the "open society" is the only way to avoid collective "submission to tribal magic" and its enforcers, "the Inquisition," "secret police," and "romanticized gangsterism."[17] Science must then be allied to this cause.

A careful reading of Popper and other antecedents of "contextual empiricism" highlight the co-productionist point that the liberal models of science are equally epistemic and ethico-political and should be argued for as such. It is not enough to say, for example, that "contextual empiricism" leads to better or more objective knowledge, as Longino tends to do. Years after *Science as Social Knowledge*, Longino implicitly acknowledges this ethico-political aspect in *Fate of Knowledge,* citing Mill, Popper, and Pierce as "predecessors" who also consider the "sociality of knowledge." Nominally, this would be a chance for her to highlight her commitment to the political vision

[16] *Open Society,* p33
[17] ibid. p195

that corresponds to her epistemology. After all, isn't Mill remembered primarily as a political and moral theorist? Longino doesn't make much of that association and restricts herself to an epistemological framing, in which "the social" is not itself normative (at least not in the philosophical sense). When she compares her view and with that of Mill, she asserts that they coincide with respect to "the necessity of critical interaction for the integrity of knowledge."[18] But wouldn't it be equally fair to say that the two theorists agree because they are committed to the idea that every human is entitled to argue for their views and be heard by diverse others? We suddenly find ourselves back in Boyle's lab, wondering where politics end and epistemology begins. Did we not already figure this out? Have we ever been modern?

"Epistemic acceptability" with a capital E

Compared to *Science as Social Knowledge*, Longino makes similar prescriptions for science in *The Fate of Knowledge*: *i)* Venues: there must be places in which critical interaction can take place, *ii)* Uptake: criticisms must be taken seriously, *iii)* Public standards: there must be shared criteria for reasoning, and *iv)* Tempered Equality: members of the epistemic community should be treated as intellectual equals. But this time, the prescriptions are framed in a different way. Instead of analyzing the concept of "objectivity," Longino claims that she is providing a social account of "epistemic acceptability." The phrasing reflects her revised project of describing the social norms that allow "effective" critical discursive interactions, the "features of an idealized epistemic community."[19] The project framing thus underplays the distinction between describing existing knowledge-productive practices and prescribing how knowledge-productive practices should work. It is a compromise that represents what Longino calls, the "non-dichotomizer's way," a rejection of the rational-social dichotomy.

I am quite sympathetic to this solution, but Longino's language is troubling in several places. Notice that the passive verb form of "idealized," in "idealized epistemic community." Idealized by whom? Ideal for everyone or just Longino? The reader is left to guess who takes responsibility for this vision and why. Similarly, the word "effective" presumes some standard of practical success, but the only standard suggested is found in Longino's choice to define "knowledge-productive

[18] FoK, p4
[19] p134

practice" as "critical discursive interaction." Without extensive empirical work showing the equivalence between these two things, there is an eerily tautological character to Longino's assertion. Mill, Popper, Peirce, not to mention Jürgen Habermas, would probably not complain about any of this, given their shared commitments to ideal discourse and liberal social order. But to everyone else, it becomes increasingly apparent that "epistemic acceptability" is just as problematic as "Truth" with a capital T; it functions as a thinly-veiled honorific term for Longino's preferred forms of sociality. The only response? To continue asking the fundamental co-productionist questions: knowledge for whom, and with what interests?

These ambiguities and omissions in *The Fate of Knowledge* contrast starkly with Longino's stance in some other works. In "Multiplying Subjects and the Diffusion of Knowledge" (1991), her very first sentence exclaims "knowledge is power." In the article she explains how traditional epistemology, Cartesian or otherwise, valorizes the individual but gives few resources to women and other persons on the margins of society. As a result, existing inequalities are amplified as those with power ignore and silence those without, often in the name of Truth. Longino asserts that her own epistemological prescriptions are an attempt to apply this feminist insight and to remedy inequality. Also in a co-productionist mode, Longino (2003) asks "Does *The Structure of Scientific Revolutions* Permit a Feminist Revolution in Science?" Her main concern there is whether Kuhn's emphasis on theory-ladenness and paradigm-dependence in scientific knowledge production unintentionally prevents us from being able to critique sexist science. Here again, it is clear that Longino is sensitive to the interdependence of our ethico-political commitments and our analyses of knowledge. So I cannot help but wonder why Longino felt it necessary to set aside these motivations in at least two of her book-length works, why feminist values appear only tangentially. What is the reader to make of these artificial separations and the total absence of "liberalism," "egalitarianism," and their potential justifications?

A short consideration of the institutional and cultural discipline of philosophy makes Longino's boundary work more understandable. While many influential epistemologists have tackled the problem of knowledge/power and taken political theory seriously, the task is somewhat different from the position of a feminist philosopher, who often faces a strange and unfair set of demands from the philosophical community. According to Phyllis Rooney (2011), feminist epistemologists have been and continue to be marginalized in favor of epistemology "proper," which is taken to be more neutral, reasoned, and apolitical. She highlights how perverse this situation is by reminding us not only that these characterizations are inaccurate but also that there is no such unified body of work that one could point to as "proper." The philosophical literature is not sufficiently homogenous to justify these distinctions. Yet, the criticisms persist. Even as recently as 2012, Kristie Dotson argues that the situation is not unique to epistemology. She suggests that professional philosophers (especially in America) are collectively possessed by a need to legitimate every intellectual project as within the scope of their field. This norm, she explains, functions to exclude diverse peoples and creative approaches. It is not a mystery, then, that Longino underplays her political commitments in her first two books. Writing in 1990 and in 2002, to an audience that is still not uniformly sympathetic to feminist projects, her language had to strike a balance between her intellectual project and the likelihood for unfair misreadings by philosophers of science.

For the purposes of my intellectual project here, I suggest we assert interpretive privilege and read Longino's epistemological prescriptions as part of a broader feminist egalitarian vision for society. Though such an explicitly political vision may not gain much traction in among many philosophers of science, I trust that it is in keeping with Longino's overall scholarly project. Despite my own hand-wringing here in this chapter, many scholars have gone ahead and done exactly this. Biddle (2009), for instance, reads Longino's epistemology as "logically embedded within the framework of Mill's political philosophy." I'm not certain about which is embedded in which, but the general idea seems right. Intemann (2011) does the same, suggesting that we take Longino's account as a "Millian" framework. She worries, however, that Longino does not fully consider the downsides of how values are represented in liberalism. One of the main problems, according to Intemann, is that Millian democracies "endorse a kind of neutrality about values that gives rise to conceptions of diversity and dissent that put racist, sexist, and creationist values on par with feminist values, as all are equally instrumentally valuable within scientific communities." Thus we run into the core weakness of liberal political frameworks; there is no principled reason to keep out objectionable value-orientations. Longino's epistemology thus may run into substantive objections of a political sort; not everyone agrees with Millian liberalism.

From whence the philosopher's authority?

A maximally charitable interpretation of Longino's framework juxtaposes her feminist egalitarianism with her "contextual empiricism" as two sides of the same coin, so to speak. But as Intemann's worries show, we must not leave this component of "critical contextual empiricism" as implicit or unquestioned, focusing only on epistemic issues. I hope that much has been established in this chapter so far. But with those matters settled, we are left with a more fundamental question: with what authority can a lone philosopher dictate a vision for knowledge in society? Longino, interpreted in the broadly co-productionist sense, proposes a compelling ideal for science, equal parts political and epistemic; her ideal society is one in which there are no unfair concentrations of power, where everyone has a voice, and where knowledge is never held fixed. It is a vision shared with the likes of Mill and Popper. But since she cannot (and hopefully does not) rely on appeals to "truth," or "objectivity," or "justice" in her arguments, she needs to explain how her own situated contribution should be weighted against competing visions, like Plato's *The Republic* or Polanyi's Republic of Science. Why should we, as fellow citizens, listen to Longino's proposals as opposed those of a different philosopher? Why listen to a philosopher at all?

Some will find my questions absurd. "Longino is a very intelligent philosopher," they might say, "so she has the relevant expertise to propose a political/epistemic ideal for society." An appeal to the philosopher's intellectual caliber or expertise is flattering to philosophers, who may be content to wield their authority wherever they can. However, a thoroughly co-productionist philosopher will not take her ability to prescribe for granted. The privilege of an elite few to envision the ideal society is tied up in historically-situated imaginaries of democracy and perhaps one of the more striking features of the present day. Cornelius Castoriadis (2010) makes this case by contrasting recent democracies with the rule of the Athenians; he argues that political judgments in Ancient democracy were always a matter of opinion (*doxa*) and within the purview of every citizen.[20] Democracy of the last century or so is imagined in a very different way. He laments in *Philosophy, Politics, Autonomy* (1991):

> The prevalent [modern] idea that there exist "experts" in politics, that is, specialists of the universal and technicians of the totality, makes a mockery of the idea of democracy: the power of the politicians is justified by the "expertise" they would alone possess, and the, inexpert by definition, populace is called upon periodically to pass judg-

[20] *Society Adrift*, p125

> ment on these "experts." It also – given the emptiness of the notion of a specialization in the universal – contains the seeds of the growing divorce between the capacity to attain power and the capacity to govern – which plagues Western societies more and more.[21]

Castoriadis does not hesitate to criticize the way "political expertise" and its claimants effectively disenfranchise the public. Fortunately, I need not pick between these imaginaries in this chapter, but the mere presence of competing imaginaries suggests that philosophers have an obligation to consider which imaginary they perform when writing books or describing well-ordered science.

Philosophers thus have a choice about what to hold firm, to leave naturalized, and what to question and deconstruct. Accordingly, it is worth considering how Longino negotiates this question of her own expertise. We can find part of an answer in *Science as Social Knowledge*, in Longino's discussion of feminist science. There, she asserts that feminist science could be read as the "neutral" option, removing bias to make science gender-free. Her own view is different. She understands her role as feminist scientist to include "the detection of limiting interpretive frameworks and the finding or construction of more appropriate frameworks."[22] Moreover, she admits that in her more direct critiques of behavioral endocrinology, she is driven by a personal preference to expand human potentiality, to increase our sense of agency, rather than allow narrow-minded science to close it down. In making these personal values explicit, Longino stresses that she takes herself to be accountable to a community outside of herself and beyond the confines of scientific institutions. As she counteracts the entrenched values in science, she is representing values somewhere outside science. It is a shame, however, that she leaves her community of "comrades" somewhat ambiguous, an anonymous crowd. Unpacking that idea in more detail would begin to answer the question of philosophical authority.

Longino's role as philosopher (rather than scientist) is somewhat clearer in the final chapter of *The Fate of Knowledge*, where she positions her own contributions in a very modest way. There, she states that philosophers should not see describe their job as "deploying zero-sum epistemologies," but instead "be sensitive to the shifting relations of multiple research traditions and the complexity of the factors that succeed in producing provisionally stable representations of nature."[23] That statement, it seems to me, implies a healthy co-productionist curiosity, but no

[21] *Philosophy, Politics, Autonomy*, p109
[22] *Science as Social*, p191
[23] *Fate*, p212

philosophical authority per se. Her own social epistemology, she mentions in passing, bears similarity with the American pragmatist tradition: "Knowledge is sought, not imprinted, and it is sought in order to achieve particular goals and is evaluated in relation to those goals. Knowledge produces the conditions of its own transformation. The growth of knowledge is not linear, but irregular, layered, and patchy." She even suggests that we think of her epistemology as "sociopragmatism," as a tentative description of how cognitive activities link up to other purposive human action. If this affinity is genuine, then it would suggest that we read her four conditions of knowledge-productive practices as contingent, open to revision as our needs change.[24]

These passages taken together hint at an implicit justification for Longino's ideal society; she is trying to represent the opinions (in the sense of *doxa*) that are more broadly shared than those of mainstream science or of a lone philosopher. But more importantly, these passages reveal the need for more comprehensive reflection on the nature of philosophical work on science, whether direct in the lab or in a scholarly monograph. I am left with more questions than answers. First, to whom am I accountable when I intervene in knowledge practices? Do I have a community in mind when I do philosophy of science? I ask scientists "who are your publics," when I should be asking myself! A second question: is it acceptable for a pragmatist philosopher to advocate for a specific vision of society? Wouldn't such an intervention amount to simply joining the fray, just bringing in one more soap box into the lab? Rather than answer these questions now, I will return to them in the final two chapters, where I present a more formal consideration of the philosopher's claim to authority.

[24] I suspect, though I do not argue for it here, that Longino's definition of knowledge productive practices applies only to a small minority of real world cases. A mismatch between more mundane pragmatic situations and Longino's conditions for objectivity may limit the applicability of her account.

References

Barnes, B. & Bloor, D. (1982). Relativism, rationalism and the sociology of knowledge. In M. Hollis & S. Lukes (Eds.), *Rationality and Relativism*. Oxford: Blackwell.

Biddle, J. B. (2009). Advocates or unencumbered selves? On the role of Mill's political liberalism in Longino's contextual empiricism. *Philosophy of Science*, *76*(5), 612–623.

Castoriadis, C. (1991). *Philosophy, Politics, Autonomy*. Oxford: Oxford University Press.

——— (2010). *A Society Adrift: Interviews and Debates 1974–1997*. New York: Fordham University Press.

Dotson, K. (2012). How is this paper philosophy? *Comparative Philosophy*, *3*(1), 3–29.

Epstein, S. (1996). *Impure Science: AIDS, Activism, and the Politics of Knowledge*. Berkeley: University of California Press.

Ezrahi, Y. (1990). *The Descent of Icarus: Science and the Transformation of Contemporary Democracy*. Cambridge MA: Harvard University Press.

Fuller, S. (1993). Science as Social Knowledge: Values and Objectivity in Scientific Inquiry. Helen Longino. *Philosophy of Science*, *60*(2), 360–362.

Intemann, K. (2011). Diversity and dissent in science: Does democracy always serve feminist aims? In H. E. Grasswick (Ed.), *Feminist Epistemology and Philosophy of Science*. Dordrecht: Springer.

Latour, B. (1990). Postmodern? no, simply amodern! steps towards an anthropology of science. *Studies in History and Philosophy of Science Part A*, *21*(1), 145–171.

Longino, H. (1990). *Science as Social Knowledge*. Princeton: Princeton University.

——— (1991). Multiplying subjects and the diffusion of power. *The Journal of Philosophy*, *88*(11), 666–674.

——— (2002). *The Fate of Knowledge*. Princeton: Princeton University Press.

——— (2003). Does the structure of scientific revolutions permit a feminist revolution in science? In T. Nickles (Ed.), *Thomas Kuhn*. Cambridge: University of Cambridge Press.

Merton, R. K. (1942). A note on science and democracy. *Journal of Legal and Political Sociology*, *1*, 115.

Mill, J. S. (1859). *On Liberty* (2nd ed.). London: John W. Parker and Son.

Popper, K. (1950). *The Open Society and Its Enemies*. Princeton: Princeton University Press.

Rooney, P. (2011). The marginalization of feminist epistemology and what that reveals about epistemology 'proper'. In H. E. Grasswick (Ed.), *Feminist Epistemology and Philosophy of Science*. Dordrecht: Springer.

Visvanathan, S. (1997). *A Carnival for Science: Essays on Science, Technology, and Development*. Oxford: Oxford University Press.

Chapter 3 - Who's Responsible in Technoscience? Philosophy of Science as Ethical Study

Ambiguous Spaces, Bivalent Practices

I'm sitting across the table from a university professor. This professor, like many others in 2015, has found a role to inhabit within the emerging field of neural engineering. And even though "neural engineer" is not yet an official title for anyone, there are many people who associate themselves with this phrase—a brand of sorts—regardless of whether they are engineers or neuroscientists or professional programmers. Federal funds, and the epistemic priorities that come with them, are often the impetus. The person sitting in front of me is no exception, overseeing a graduate student project on neural interfaces. The National Science Foundation facilitates these interactions by creating research clusters, like the Center for Sensorimotor Neural Engineering (CSNE), and funding primary investigators across the country for one extra student. I ask about what neural engineering means and where it falls among the existing categories of science and engineering:

> Interviewer: What does "neural engineering mean to you? ...just a rough description.
> Interviewee: It means a lot of things I guess. As an educator, it's an opportunity to teach people about this interesting new interdisciplinary field that combines science and engineering. It has medical applications. So there's that aspect; it's a great vehicle for education. And then on a research level, it is a way of helping people.
> Interviewer: And when you describe your research, do you describe it as engineering research or scientific research or maybe neither?
> Interviewee: It depends on the audience. For some people, I don't necessarily make a distinction, and I just use "science" because it's easier. People are more attuned to it. But for people who do understand I will call it engineering. Unless sometimes we delve into science, but usually we are builders of stuff.[1]

As this exchange illustrates, discourse around neural engineering tends to emphasize the interdisciplinary character of the field or blur the distinction between "science" and "engineering" research. The NSF, requesting proposals for third generation engineering research centers, describes their goal as creating "an interdisciplinary research program that aligns systems-motivated

[1] This and other dialogues come from a set of interviews conducted at the Center for Sensorimotor Neural Engineering. See Chapter 4 for details.

fundamental and applied research with enabling and systems technologies."[2] It's not evident from this document how systems-motivated research can be fundamental.

Confused, I asked a CSNE-affiliated neuroscientist, working on a similar project, about disciplinary identity:

> Interviewer: Do you understand yourself as primarily as an engineer as opposed to a scientist or something else?
> Interviewee: I primarily understand myself as a scientist who is involved in the generation of tools to ask sophisticated questions with respect to "how does the brain work?" and "how can we restore it if it doesn't?"
> Interviewer: And would you say that you're conducting scientific research or engineering research, or...?
> Interviewee: I would say I am conducting scientific research.

Thus we have two people, working on very similar projects with similar goals, who define themselves as "scientist," as "engineer," or both. There is a flexibility available and exploited within these researchers' language, a discursive strategy that can also be seen in the way that the "basic research" is regarded in the CSNE.

In 1945, Vannevar Bush famously called for federal support of "basic research" which is conducted "without thought of practical ends" and produces understanding of "nature and its laws."[3] This proposal, with its vision of a National Research Foundation, reimagined the role of scientists in the United States, ostensibly freeing some from the obligation to develop military technologies and creating a space for relatively autonomous research. And yet, the idea of "basic" research did not fully take hold in the resulting institutional structures. The NSF forbids such pursuits in ERCs: "Basic research activities without clearly presented strategies to advance enabling and systems technologies and proofs of concept will not be supported."[4] The idea of the ERC's "fundamental" research both draws on and contradicts Bush's call for descriptive, curiosity-driven science.

The CSNE, in particular, fulfills this ambivalent mandate via the production of "principles" of brain-computer interfaces. This output is neither simply empirical knowledge of the natural world, as one might expect from "basic" research, or functional proof-of-concept devices, as engineers often create. If you ask researchers to explain, they often link the principles to engineering, but are happy to describe the more scientific elements inherent in the research:

[2] NSF ERC Gen 3 solicitation, 2013
[3] Bush (1960) *The Endless Frontier*
[4] Gen 3 ERC solicitation

> Interviewer: The CSNE recently defined its output in terms of principles. In your understanding, what sort of principles are these? What does "principles" mean in this case?
> Interviewee: The way I think about principles is as guidelines as opposed to best practices or you know... It depends on which aspects. Pulling from the principles most closely related to what we do...principles of long term stable neural interfaces are: "what do you need?" ...a list of features and qualities that are important or not important or essential in chronic long-term interfaces.
> Interviewer: For the CSNE principles, the ones your lab comes up with, would you say these principles are better understood as representing features of the natural world or as defining feasible human interventions? Or neither? Does that distinction make sense?
> Interviewee: I think they're about engineering. So I consider them to be whatever that other one was.
> Interviewer: Defining human interventions or...
> Interviewee: Ya.
> Interviewer: ...the rules of application or something like that.
> Interviewee: Ya.
> Interviewer: And would you agree with the statement that there might be moments of basic science in producing these principles?
> Interviewee: Oh sure. Oh ya. I mean as you, for instance, uncover the principles for chronic interface we need to understand something—we don't need to, but it's probably helpful—to understand how the body is reacting to these. And *that*...is science.

The interviewee, here, is citing the role of histology at the CSNE; in the process of putting new materials in the human body or in animal models, they learn new things about cellular and immunological response. Whether or not these findings are eventually instrumental for the project, they have a descriptive component. Cell biologists at the CSNE employ standard biological and biomedical measures of inflammation and follow their discipline's standards for robust microscopy. But that "science" is demand-driven, so to speak, unique and beholden to the neural engineering context. No one outside the field is implanting polydimethylsiloxane structures into animal brains, and discovery or curiosity are secondary concerns, at most.

Looking more closely at the institutional structure of the CSNE, we find an organizational analog to the blurred identities and nested practices described above. The CSNE unites biologists, neuroscientists, electrical engineers, surgeons, and others under one set of goals—funding is contingent on contributing to shared goals—and one physical roof in Seattle, WA. The CSNE has a dedicated "industry advisory board," an "practitioner and end-user advisory board", and a "scientific advisory board" (see fig. 1). As indicated in CSNE depictions of its structure, the ERC leadership consistently promote an awareness among members of their complex relationships with various administrators, stakeholders, and interests. Even the public-facing web presence of the CSNE draws our attention to these institutional arrangements, carefully reconstructed online for whomever happens to visit their page.

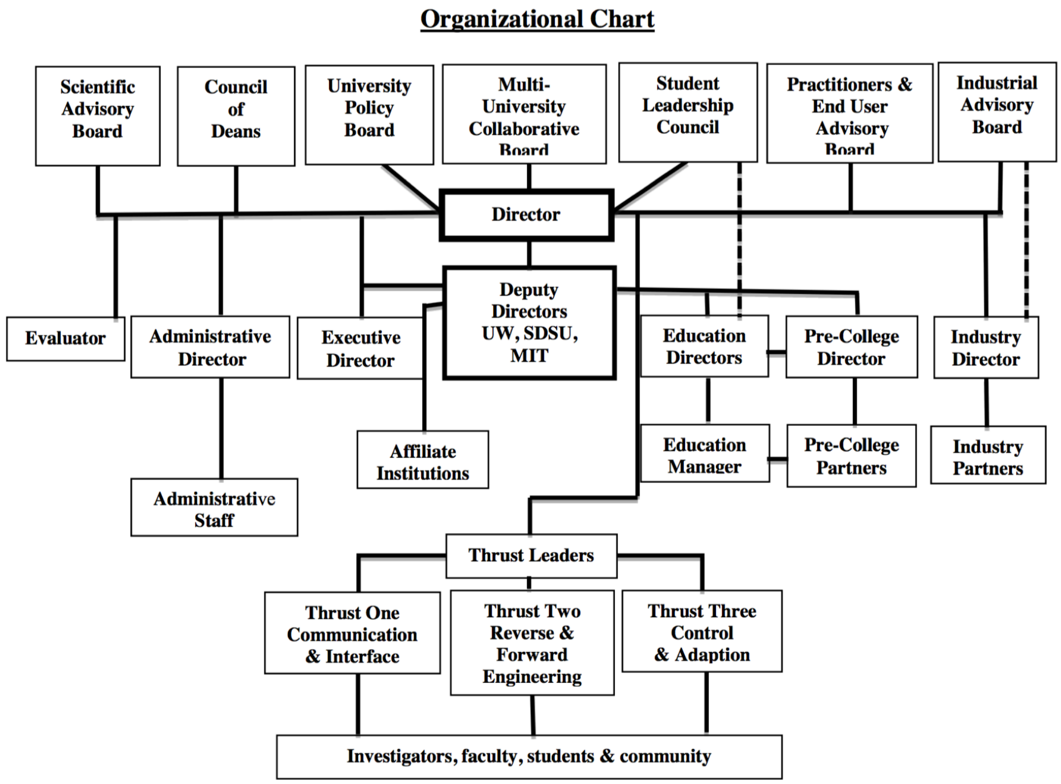

Figure 1: CSNE organizational chart, as presented in a 2012 internal report

Taken together, the language, self-understandings, and institutional self-representations presented briefly here begin to reveal a parallel failure in the common disciplinary and philosophical distinctions between "science" and "engineering", "basic" and "applied." These terms can't fully describe contexts like the CSNE.[5] And yet, the writers of NSF solicitations and actors within the CSNE seem to have little difficulty working around the conceptual ambiguity. As the interviewee above pointed out, the official description of CSNE research depends on the audience. In the same way, philosophical distinctions between ethical, political, and epistemic do not always pick

[5]There is reason to suspect that they never fully *described* anything, if we look at their political and social functions.

out separate activities in technoscientific contexts. The "fundamental research" or "science" at the CSNE, if we can identify at such, is mission specific; it is conducted in anticipation of future translation into neural technology. The translational project, similarly, is not a matter of merely creating neutral tools for society or leaving the final purpose of the technology open-ended. As I described in the introduction, the push for neural technology is joined to a political vision in which people with disabilities voluntarily and surgically "restore" themselves.

While an epistemologist could choose to ignore these features, a practice-sensitive philosopher of science cannot. Faced with the hybrid character of modern technoscience and the multivalent life of researchers, one promising approach is to focus on responsibility. By grouping the "oughts" of knowledge creation with the "oughts" of our broader moral lives, we are given a way to re-integrate epistemology into a more socially sensitive research paradigm. In this way, we might fulfill the vision I lay out in Chapter 1 and acknowledge the co-production of technoscience and society. Specifically, the responsibilist approach is sensitive to validity claims beyond the solely epistemic. It might have the potential to deal with political trajectory of neural technologies, adding justice to the epistemologist's toolkit. Thus, by guiding philosophers of science onto the path walked by ethicists and moral philosophers, a responsibilist lens seems a promising answer for the co-productionist challenge for philosophy of science.

In this chapter, I will explore the origins of the responsibilist approach and present the work of Heather Douglas as an attempt to apply the framework to technoscience. Along the way, I will use a few examples from technoscientific practice to show that a responsibilist focus is in tension with the lived, social experience of many scientists and engineers. By juxtaposing responsibilist descriptions of science with self-descriptions from within neural engineering, I suggest that there are at least two contradictory imaginaries: the speculative philosophical picture of science and the self-understanding of neural engineers. The mere fact of this mismatch is intellectually significant. However, if we desire to fulfill the responsibilist program—it is meant to be applied after all—then we must find ways to reconcile these two understandings of one social world. I will conclude the chapter by briefly considering the feasibility of this reconciliation.

The Moral Life of Knowledge

In "Ethics of Belief" (1877), Clifford makes a bold suggestion. What we believe can be as ethically significant as what we do. We are obligated, he says, to believe only according to the evidence at hand. To violate this obligation is not just an epistemic failure; it is a moral one. While Clifford's dicta in that essay are perhaps too general and too strict to be useful for a consideration of technoscience, the idea of an ethics of belief is crucial to the responsibilist approach. Through this lens, we can explore the possibility that doxastic obligation and other forms of obligation might originate from the same place, from the norms of community membership. We can thus begin to explore the commonalities between true belief and right action.

Within philosophy, it is feminist scholarship that has taken this idea to its logical endpoint, expanding the scope of ethics. Lorraine Code (1984), more recently, gives epistemology its "responsibilist" formulation. Navigating a middle path between foundationalism and coherentism, Code asserts that there are important analogies between the ways that ethical judgments are justified and the ways that we take something to be knowledge. She proposes that epistemology should be rebuilt in terms of "intellectual virtue." This approach has three important characteristics. First, Code identifies epistemic responsibility as the primary intellectual virtue, similar to but not co-extensive with Aristotelian virtues like courage or prudence. Second, intellectual virtue can be construed as a sort of sensitivity to the world and one's place within it. Finally, the virtuous knower can take different forms, depending on her nature, her environment, and—particularly important for my project—her epistemic community.

The effect of this virtue-based account is to shift our attention away from the knowledge status of particular beliefs and towards a person's character and social situation. In this way, we can make our epistemology "social," but without reducing the individual to a cog in the knowledge machine. There is no need to multiply abstract knowers using computer-based models and artificial networks. Instead, we can begin to examine the rich communal structure of knowing and still treat the individual knower as an agent, who might be worthy of praise or blame.

Sometimes, the knower's position within a community form of life entails additional responsibilities in virtue of her role in society. Someone who takes on the role of math teacher, for instance, is expected to know algebra quite well, unlike a philosopher. Other times, the contingent character of intellectual virtue dictates that we withhold attributions of responsibility. A color-blind person, for example, might be labeled as having a deficiency in visual senses, but he could not easily be held responsible for confusing some shade of red from some shade of green. Similarly, Code suggests, we could not label a Soviet scientist as irresponsible for not reading Western science; it would have likely been inaccessible due to the structure and policies of Soviet scientific institutions.

Scientific inquiry, as a knowledge-oriented practice, fits within this broader epistemological framework of intellectual virtue. Drawing on Polanyi (1966), Code (1987) suggests that it is the communal commitment to realism that shapes epistemic responsibility within science; "In scientific practice at its best, then, it is normative realism is the apparent force that guides and shapes the progress of inquiry." But how does this shaping affect the individual? A scientist's particular place within the community, a role which she plays, defines the duties and limits of epistemic behavior. So a successful state of knowing, one might say, is a function of the epistemic norms imposed by one's community. "Normative realism" stipulates that the scientist can only engage in experiments and inferences that takes some part of the world as given, waiting to be described. This norm entails, for example, that the scientist cannot justify holding onto a belief only because she wants it to be true of the world.

Applying Responsibilism to Technoscience

In *Epistemic Responsibility*, Code begins to spell out a perspective within which we can hold technoscientific actors accountable for their stated beliefs and their knowledge claims. There is a definite meaning to "responsible" knowing within science: accordance with normative realism, balancing personal thought with institutional wisdom, a humble fallibilism, and so on. This makes science a "practice" in the usage of Alasdair Macintyre; it is a space in which we can exhibit virtues. Unfortunately for the purposes of this project, Code's *Epistemic Responsibility* does not prescribe what that accountability should look like within a broader community, that of democratic society. Her reliance on Polanyi entails that science is presented as a more or less isolated practice, subject only to its internal communal norms. Or as she quotes from Polanyi, Code suggests we envision an interdependent community, distributed, with "chains of overlapping neighborhoods extending over the entire range of science."[6]

This investigation of an isolated practice raises an implicit barrier between society and science. In this separation, we see the traditional philosophical vision of science guiding analysis, while the co-production of science and society is set aside. Our attitudes towards technoscience and the obligations we impose upon it, are thus left unresolved. This consequence is not lost on Code. She juxtaposes Polanyi with Michel Foucault in order to show how discourse is not a politically-neutral entity. It is up to members of society, especially those within institutions of knowledge-making, to cultivate a sensitivity to the connections between expert discourses and the exercise of violent power. "Thinking individuals," she argues, "have a responsibility to monitor and watch over shifts in, changes in, and efforts to preserve good intellectual practice."[7] And as one would expect of epistemic *virtue*, good knowing is defined as a balance between fulfilling ones community obligations and resisting violent power-knowledge structures. The good knower must not let collective expectations dictate dysfunctional forms of knowledge, but neither can she neglect her practice and forget her peers in the community of inquiry.

What could Code's middle-way look like for technoscientific researchers in the present day, if we expand the community of science out to include the public? Heather Douglas' recent work addresses this question and begins to close this gap between ideal and non-ideal science in her work on the obligations of scientists. Her arguments in *Science, Policy, and the Value-Free Ideal* (2009) continue a conversation that dates back to the mid-20th century, an unfinished debate over whether scientists deserve "autonomy" in their pursuit of knowledge. As Douglas describes it, one side—Michael Polanyi is a central figure here too—emphasizes the ways in which science is and must be free of social or political values. These arguments for autonomy, relying on "tacit knowledge" and scientific realism, represent several compelling possibilities. But value-freedom, even as an ideal, is not quite so simple. Douglas reminds us of the other side of the debate. Responses from Rudner (1953) and Churchman (1948) reveal that we cannot avoid judging the quality of a particular hypothesis or the risk we take on by accepting it.

With these historical points and counterpoints laid out, Douglas presents her own analysis. Scientists, she suggests, have "general" moral responsibilities; they are responsible for their behavior in the way all social beings are. And they have "role responsibilities," unique to the position of a

[6] Polanyi in *The Tacit Dimension*
[7] *Epistemic Responsibility*, p245

scientist. She acknowledges that, sometimes, our roles can take priority and remove our general moral liability. Lawyers, she suggests, have a role that shifts responsibility onto the legal system and can disregard some general moral responsibility as they defend a guilty client. So we might think that the role of scientist is like that of a lawyer, such that general moral responsibility for mistakes and oversight is set aside. Douglas thinks not.

The problem, she suggests, is that usually there is no "rigid" institutional structure within science that can take morally significant decisions out of the hands of scientists.[8] Unlike the context of law, there is no adversarial system to ensure diverse or antagonistic perspectives are marshaled. There is no judge that at the end of the day has the power to decide whether or not a hypothesis deserves further testing, weighing its probability against the very real costs of a false positive. This is a decision that only the scientist and her immediate colleagues can make. Douglas, thus, argues that scientists have an obligation to consider the cost of error in their professional claims; they are not off the hook when it comes to general moral responsibilities.

In virtue of its responsibilist perspective, Douglas' analysis seems like a promising extension of Code's epistemological program. It may provide us with a tool with which to normatively assess technoscience, capturing both epistemic and socio-ethical dimensions of concern. As described in *The Moral Terrain of Science* (2013), Douglas presents an expanded vision of science; it is a practice that is situated in and valued by society. She uses this fact to prescribe scientists' responsibilities beyond internal communal norms. There are, accordingly, three "bases" of responsibility to which scientists are accountable: "The first is to good reasoning practices; the second to the epistemic community of science; and the third to the broader society in which science functions and is valued." If we want to normatively evaluate a particular research project or field—recall the CSNE's potentially controversial research on neural interfaces—we can run through each basis.

On the first basis: did the research produce "reliable empirical knowledge"? According to Douglas, science exists because society values it as a source of "reliable empirical knowledge," and is thus beholden to that epistemic standard.[9] *On the second basis:* did the researchers support and enable their epistemic peers? As members of an interdependent community, scientists cannot act as if they work alone. *Finally, on the third basis:* Can we say that neural interface research embodies and advances the values attributed to science by broader society? Since science does

[8] SPVFI, p73
[9] ibid. p95

not occur in a vacuum, research and its reasonably foreseeable effects should not conflict with societal values.

As Douglas points out, this framework prevents anyone from problematically bracketing questions of knowledge from questions of morality; scientists can't excuse themselves from moral responsibility (i.e. basis 3) for the sake of empirical knowledge (i.e. basis 1). In this way, Douglas' framework might seem well-equipped to deal with the potentially controversial research at the CSNE. A neuroscientist, for example, cannot cite her dedication to good science simply to evade normative evaluation of her contribution to technological development, whether for military lie-detectors or novel treatments for depression. A responsibilist philosopher of science is likewise encouraged to develop a sensitivity to the whole range of obligations that scientists might have, whether epistemic or ethical or in between.

Figure 2: Douglas' idealization (imaginary) of science

This strategy seems argumentatively effective against those who espouse the "value-free" ideal of science. However, Douglas' responsibilist approach still draws on her own philosophical imaginary of science, rather than an actual existing practice. Just as Code's vision of science in *Epistemic Responsibility* draws on Michael Polanyi, Douglas' descriptions of science seem to continue Kuhn's legacy of treating science as self-contained community, which operates according to internal norms. This choice is most evident in her non-institutional definition of science: "an iterative, ampliative process of developing explanations of empirical phenomena, using the explanations to produce predictions or further implications, and testing those predictions. In light of the new evidence from the tests, explanations are refined, altered, or further utilized." Looking at minimal diagram of her vision (Figure 2), it is not clear if the scientific community (presumably an institution or an organization rather than a process) would be correctly circumscribed by such a definition; if we took the process she describes as the definitive norm for scientists, then the "scientific community" would exclude much of what I see at the CSNE and (as I will describe in the next section) what STS scholars highlight as "technoscience." It also fails to account for clinicians, DARPA, industry, the Market, bioethicists, scientist-engineers, and other things that are not self-evidently inside or outside the "scientific community" thus defined.[10]

If our goal is to use idealizations to improve society, then Douglas treatment of science is unsatisfactory in a couple of ways. First, an ideal with little connection to practice may inhibit its utility. If a philosopher—Code or Douglas included—uses a non-existent or ideal practice to set its value and its community structure, the resulting prescriptions may be inapplicable or practically useless. The prescribed obligations will have little or no force for technoscientific researchers, who may not recognize the imaginary being presupposed. "Who said we're all in the business of explaining?" they might say. Second, any sociotechnical imaginary comes with assumptions about how the world ought to be. Even if Douglas admits the loftiness of her conception of science—maybe it's just a distant hope—then there remains the task of justifying that desirable state of affairs to persons who might not share Douglas' value judgments. Returning to some empirical findings from neural engineering practice will help make these criticisms more concrete.

Philosophical Imagination vs. Technoscientific Self-Understanding

We should primarily understand Douglas as responding to a particular debate about value-free science. In that sense, she presents a necessary correction to the ideal of science as value-free by formulating an opposing ideal. But if her mapping of the "moral terrain" is to provide a new responsibilist foundation for philosophy of science, there are some additional obstacles to consider. In Douglas' main arguments, she treats science as a single identifiable practice, independent from university professorships, federal administration, and "neuroventure" start-ups. Recall, for example, that her imagined scientific community is circumscribed by its pursuit of "re-

[10] Refer to Figure 3 in Chapter 4 to see an alternative idealization of technoscience.

liable empirical knowledge." Douglas also distinguishes between "the value of knowledge" and "social value," which (without some qualification) neglects her overarching insight that scientists frequently make policy and, in general, influence social order. This is disappointing given that Douglas herself acknowledges the extent to which philosophy of science, under the influence of Thomas Kuhn, has artificially bounded science.[11] On one reading, Kuhn uses *The Structure of Scientific Revolutions* to hide a vision of autonomous science within socially-rich imagery, to create a "social" picture of science that neglects society. Steve Fuller, for instance, sees this as *the* consequential and anti-democratic misstep within philosophy of science.[12]

According to some science and technology (STS) scholars, there are concrete reasons to avoid Kuhnian boundary-drawing. Alfred Nordmann (2012), for example, stresses that much of contemporary science ("technoscience") is oriented towards "knowledge of control" and capacity-building. This finding is borne out in my own qualitative research on neural engineering practice. Remember that much of the research at the CSNE is labeled not as "basic" but rather as "fundamental" and "systems-motivated." The research is not justified to 'the public", generally speaking, but instead to advisor-stakeholders in industry, to surgeon-collaborators and members of the disability community. The action of institutional structure, funding priorities, and disciplinary identities, among other things, complicate Douglas' partially *a priori*, partially empirical attributions of scientific responsibility. It is helpful here to consider how these complexities are understood by the individual researcher in the CSNE; the researchers sketch out an imaginary that looks radically different from Douglas'.

The social world of the neural engineering researcher is primarily the subject of the next chapter, where I give a full analysis of interview data. But to demonstrate the contingent character of Douglas' picture of science, it is sufficient to consider a representative interview. I ask one professor about who determines the "direction of research":

> Interviewer: Are you in full control of the direction of research in your lab? Maybe not just [your student's] project, but also in your lab in general...
> Interviewee: Full control? I hope not. I would like the students to think that they have an opportunity to be inquisitive on their own and go in their own direction as opposed to being told what to do at all times. I mean, after all, we're trying to educate students who are independent thinkers so that's kind of "anti" to being in full control. There's

[11] SPVFI, p61
[12] Fuller (1992) "Being There with Thomas Kuhn"

> some issues that I better be in full control of like safety, monetary responsibility, and stuff like that.
> Interviewer: Okay. But as far as the actual direction of particular projects, you give a lot of autonomy to the student?
> Interviewee: Well, we have deliverables and our sponsors expect us to get the deliverables, but, you know, while we're meeting deliverables, I would hope the students feel that they have an ability to go their own way at times.

In this short interaction, we find traces of a complex network of responsibility relations. The professor monitors financial and safety concerns at the lab level, while meeting pedagogical ideals. Sponsors set the terms of research success. And the student is, in some sense, free to exercise judgment within those parameters. In the future, once neural technologies exist, will we remember this complexity in their origins, or will we have in our heads the vision of a lone engineer crafting a brain-computer interface? Future critique of or enthusiasm for neural technology may lead us to imagine some site—or person—on which to locate praise or blame. This hypothetical draws out the difficulty inherent in the idea of "responsible technoscience"; responsibility tends to attach to agents (collective or individual), but in the eyes of its practitioners, neural engineering seems to be a distributed phenomenon.

Previous work in STS has shown that researchers will often describe themselves as cogs in a machine, lacking true control—and thus responsibility—over the direction of their projects.[13] Directed back at Douglas' "bases of responsibility", this realization can be used to pose a very simple critique of her prescriptive claims. An individual can't be held responsible without some combination of freedom, causal sufficiency, and factual understanding of possible consequences. And technoscientific actors might deny having any of the three conditions! If, for example, funding agencies only further unethical values, then the individual researcher is faced with a severe constraint. Or similarly if the disciplinary identity of a neuroscientist precludes ethical reasoning, than the failure may be placed on "scientific culture" rather than on scientists.

This critique, however, is overly simple for a couple of reasons. Douglas acknowledges that individuals may have to "collectivize" certain responsibilities that are too big for one person, as in the creation of the National Science Foundation or the National Institutes of Health;[14] a lack of causal sufficiency is not an automatic excuse for the individual. Thus, individuals may have a second-

[13] see for example Swierstra & Jelsma (2006)
[14] Moral Terrain, section 4

order responsibility to work around and/or change institutional obstacles to their fulfilling their obligation to the "three bases." Moreover, she acknowledges that we need to recognize the range of between "minimum standards" that are not "absurdly strict" and "ideal behavior" that may be unattainable for anyone. Douglas, thus, acknowledges the existence of social constraints on individuals and attempts to integrate that insight into her prescriptions.

Yet, in meeting the simple critique, Douglas commits philosophy of science to a profoundly sociological and anthropological task. Responsibilist philosophers of science, if they are to genuinely evaluate technoscientific practice will be required to completely reshape their traditional conceptions of science, integrating all the messy cultural and social details of each case. They have to acknowledge the existence of advisory boards and the researchers' obligations to the "end user." Consider, for example, the analogous requirement in Code's responsibilist epistemology: "As theorists of knowledge, we need to be reasonable in our expectations in order not to impede genuine possibilities of insight with the imposition of unattainable goals. Epistemic responsibility is a stringent requirement, but not an impossible one.[15]" Here, Code points out that intellectual virtue, as a matter of character, has to take the constraints of reality into account. We must, therefore, make our expectations compatible with the particularities of humans who have limited physical bodies and are part of a particular community, which may not be neatly organized around "empirical reasoning."

Correspondingly, Douglas' "three bases of responsibility" must be "stringent" but not "impossible" with respect to the context in which actors exist. This distinction requires additional clarification. Saddling technoscientific actors with "stringent" responsibilities entails a sensitivity to their position in the world. In my personal experience, however, a philosophical education does not provide any such sensitivity. I was oblivious to the actual practice of technoscience until I started working near labs and applying for federal funding. Accordingly, philosophers should double check that their idealizations of science are not completely detached from practice. To return to the observations that opened this chapter, we may not be able to write books about responsibilities of "scientists", when that label has expanded to cover engineers and businesspersons as well. *Science, Policy, and the Value-Free Ideal* may need to be expanded and revised to *Technoscience, Policy, and the Value-Free Ideal*. The "three bases" will probably take on entirely

[15]Code (1984), p50

new meanings when the "epistemic community" is mapped onto the institutional arrangements depicted in Figure 1 (and Figure 4 in the next Chapter).

Moral Philosophy (of Science)

In sum, philosophers of science must draw on new, perhaps unfamiliar, scholarly methods if they are to think about technoscientific responsibility in a rigorous and reflective way. Sociological work on structure, for example, can organize our thoughts about how the individual navigates her material and social environment. And not just any sociology will do; we need a sociological framework that strikes a balance between treating actors as knowledgeable sources of social insight, on the one hand, and as politically-vested promoters of particular visions of science, on the other. Here, a combination of Giddens (1984) concept of "structuration" and the anthropological concept of "imaginary" fits the bill. In the next chapter, I will draw on these theoretical resources in order to interpret a set of interview data and capture the community structures and technoscientific imaginaries relevant to neural engineering.

Assuming that we can create an adequately detailed understanding of structuration in neural engineering, how can we use it to inform our expectations of responsibility? This question can be broken down into several others. Who gets to set reasonable standards of responsibility? The answer should probably not be, "just philosophers." Negotiating the relationship between technoscientific practitioners, their institutions, and various publics is a task for broad and democratic deliberation. Along similar lines, we can ask how are the "reasonable standards" of responsibility should be chosen? The outcome would be very different if we proceeded via principles instead of procedurally.

In the end, it may be impossible to attribute responsibility from a narrowly philosophical perspective. But this isn't necessarily a reason to give up on the responsibilist framework or on philosophical scholarship. Doorn (2010), for example, suggests that the only way to come up with relevant ethical principles is to produce them procedurally with all actors involved. In that way, the resulting requirements could be "stringent" but would reflect the actual lived experiences and social worlds of actors. Or alternatively, philosophers could collaborate with their peers in sociology, anthropology, and political science to develop a more nuanced picture of how science is valued by the public(s) and how it enrolls industry or the state. This interdisciplinary work could be interpreted as a fulfillment of Douglas' philosophical program, remaking the "bases of responsibility" according to a more socially-sensitive conception of technoscience. In either case, it's bad news for the armchair philosopher of science or living room bioethicist. The hybridity of contemporary science and engineering forces us out of the office and into the field.

References

Bush, V. (1945/1960). *Science, The Endless Frontier*. National Science Foundation.

Churchman, C. W. (1948). Statistics, pragmatics, induction. *Philosophy of Science*, *15*(3), 249–268.

Clifford, W. K. (1877). The ethics of belief. In R. Madigan (Ed.), *The Ethics of Belief and Other Essays*. Amherst: Prometheus.

Code, L. (1984). Toward a responsibilist epistemology. *Philosophy and Phenomenological Research*, *45*(1), 29–50.

——— (1987). *Epistemic Responsibility*. Hanover: University Press of New England.

Doorn, N. (2010). A Rawlsian approach to distribute responsibilities in networks. *Science and Engineering Ethics*, *16*(2), 221–249.

Douglas, H. (2009). *Science, Policy, and the Value-Free Ideal*. Pittsburgh: University of Pittsburgh.

——— (2013). The Moral Terrain of Science. *Erkenntnis*.

Fuller, S. (1992). Being there with Thomas Kuhn: A parable for postmodern times. *History and Theory*, *31*(3), 241–275.

Nordmann, A. (2012). Object lessons: towards an epistemology of technoscience. *Scientiae Studia*, *10*(spe), 11–31.

Polanyi, M. (1966). *The Tacit Dimension*. New York: Doubleday.

Rudner, R. (1953). The scientist qua scientist makes value judgments. *Philosophy of Science*, *20*(1), 1–6.

Swierstra, T. & Jelsma, J. (2006). Responsibility without Moralism in Technoscientific Design Practice. *Science, Technology & Human Values*, *31*(3), 309–332.

Chapter 4 - Between Structure and Imaginaries: The Social World of the Neural Engineer

In 2012, I joined the “Neuroethics Thrust” at the CSNE to conduct a comprehensive socio-ethical intervention. We sat down, one on one, with Center project directors—there were almost 30 in total—and asked each of them about the ethical issues in their research. Once the interviewee had responded, the interviewer would politely but assertively suggest issues that the interviewee should be considering but somehow missed, at least from the perspective of applied ethics. The ambitious goal for this first effort was to provoke reflection on the part of CSNE researchers, to help them develop a sensitivity to the ways in which their day-to-day labwork had socio-ethical significance. The questions were often leading, and the tone was often Socratic: “By restoring ‘normal’ function, aren’t you defining ‘normal’?” While we were satisfied with this intervention as such, our aggressive approach did not lend itself to understanding the perspectives and self-understandings of CSNE researchers. In terms of qualitative research, there were no easily usable data, no transcripts, notes, or audio recordings.

The experience was instructive, nonetheless. I noticed in that first project that there was a shockingly poor fit between the language with which we philosophers spoke and the practice described by researchers in neural engineering. There was, at one extreme, individualist ethical interrogation about “issues in *your* research” or things for which “you are responsible” and, at the other extreme, interviewee descriptions of technoscience that almost lose the individual researcher entirely. Their speech directed my attention to clinics, industry, and funding agencies among other things. My interactions with Center scientists and engineers slowly began to convince me that our perspectives as “outsiders”—outsiders trained to see philosophically—challenged the insider perspective that sees technoscience as distributed, unpredictable, and collective. In this way, I found myself faced with a conceptual puzzle. How can we make sense of interaction between the technoscientific researcher, her agency, and the increasingly evident effects of social structure? What theoretical tools are appropriate? This chapter makes up the substantial answer to these questions, presenting Giddens’ “structuration” and the STS concept of imaginaries as a potential solution.

The methods and findings presented here also dovetail with the conclusions of the previous chapter. There, I conclude that advancing Douglas' responsibilist framework requires an adequate understanding of the social spaces of technoscience; philosophical imagination, as we've seen, is not always enough. But where to start? What would it look like for a philosopher to conduct sociological investigations, to be in the world rather than just think about it? The findings presented in this chapter spell out some concrete answers to these questions, but must be grounded in the conditions of their creation. As a graduate member of the "Neuroethics Thrust" at the CSNE, I have had privileged access to the spaces and people of neural engineering in the United States. Yet, my position and identity as "embedded ethicist" within the CSNE comes with its own valence of meanings, which I cannot completely control or externalize in my interactions with other CSNE members. Here in Chapter 4, I try to take advantage of this access to the social world of technoscience while acknowledging my situatedness. I attempt to elucidate some relevant social structures of technoscience while accounting for the extremely specific relationship between the those of us in the Neuroethics Thrust and our non-philosopher colleagues in the neural engineering community. Structurally, the first half of this chapter is dedicated to social theory and qualitative methods. I lay out the key scholars and frameworks that enabled me to make sense of my experiences and give them a conceptual interpretation. In short, I present the "theory/methods package" that enables a philosopher to do fieldwork.[1]

In the latter half of the chapter, I hope to draw out several specific points suggested by my qualitative findings, each with its own consequences for responsibilist and other philosophical analyses of technoscience. First, technoscientific actors consistently portray themselves as constantly traversing and ignoring boundaries (e.g. between technoscience and the public, between science and engineering, etc). Second, the interviewees oscillate between claiming to help persons with disability and admitting the inadequacy of their technological toolkit. Third, technoscience is presented as a hopelessly distributed effort in a way that inhibits easy attributions of responsibility. All told, these structures are best understood not as objective facts of technoscientific practice but as elements in the imaginary that drives and sustains efforts within neural engineering. The neural imaginary, as revealed here, is both a reflection of actual technoscientific practice and a direct response to "the ethicist in the room," namely, me.

[1] The concept of a theory/method is described in Clarke & Star (2008)

High theory looming: social ideas for the confused philosopher

Given my experiences interacting with PIs, faced with the incompatibility of individualist philosophical perspectives and the researchers' sense of collective practice, I needed tools. I sought out any theoretical frameworks that could orient the individual actor within a matrix of structural conditions and, more fundamentally, give me a vocabulary to talk about what I was hearing. My own philosophers' toolkit did not tell me what I could gain from coding and analyzing interview transcripts or how to manage my own outward identity as an embedded philosopher. I found this absence problematic; even though it seems obvious that an ethicist with an audio recorder can learn *something* from interviewing researchers, the interview dynamic may provoke particular responses from interviewees. Some of us in the Neuroethics Thrust were deeply "embedded," welcomed into labs as honorary members or "resident philosophers" who are expected to somehow make the research better. Other times, my emails requesting lab access were left unanswered for many months. These dynamics are part of the data in this chapter. More generally, any interview has to be interpreted with care. Interviewee responses are just that, responses, and not some revealed mental state or neutral fact ready to be tallied. As I inquired into the ethical concerns, social lives, and research protocols of interviewees, these actors were in a sense justifying themselves to the Neuroethics thrust, as a sort of proxy of the public.

The interviews analyzed in this chapter, thus, combine insights into the social lives of neural engineers and their strategic self-justifications to an inquiring ethicist. The data are both descriptively rich and laden with normative content. My strategy for dealing with this tension was to rely on two loci of social theory: Chicago School symbolic interactionism, as approximated by Howard Becker and Anthony Giddens, and the STS framework of imaginaries. Between the two, I treat my interviewees as active responders to their social environment, as social beings who cognitively or conceptually define the situations to which they are responding. To make this idea clearer in what follows, I will focus on two questions: 1) why should we listen to scientists and engineers? and 2) can we trust what they say to be "unbiased"?

Why listen to social actors? Giddens and structuration

I've been asked by fellow philosophers and even some STS scholars,"why listen to scientists?" and "why does it matter what they think about their responsibilities or lack thereof?" There are a diversity of reasons to do so. Harold Garfinkel's concept "ethnomethodology" stipulates a methodological priority of the actor's perspective as opposed to the inquiring social theorist.[2] Drawing on standpoint theory, Dorothy Smith (1987) calls for a feminist sociology that begins with the actor's perspective rather than the generalized relations of social order. These theorists, among many others, would have been excellent starting points and justifications for my own qualitative work. But as it turned out Giddens' concept of "structuration" was the first and most parsimonious explanation of my hunch, namely, that responsibilities are closely tied to the social structures in which we participate. Giddens (1984) provided me with some much needed sophistication in my thinking about ontological status of "structure," its interaction with social actors, and the ways in which we might learn about it.

Fundamental for Giddens' social theory is the "duality of structure." He describes the duality as a "theorem" that asserts:

> The structural properties of social systems are both medium and outcome of the practices they recursively organize. Structure is not 'external' to individuals: as memory traces, and as instantiated in social practices, it is in a certain sense more 'internal' than exterior to their activities in a Durkheimian sense. Structure is not to be equated with constraint but is always both constraining and enabling.[3]

Structure, in other words, is both the basis for human action and the result of human action, which constantly recreates structure. These reciprocal interactions can be understood as "structuration.[4] Giddens is careful to highlight the active connotation of "structuration," encouraging us to see structure not as a static state of affairs but rather a site of on-going human work.

On the agent side of the duality, we have the "knowledgeable agent." This idea sharply contradicts the picture of humans as cultural dupes or unthinking automatons following invisible social laws. Giddens thus departs from Durkheim and other social theorists, denying that actors are just points pushed about by a social field. Agents are "knowledgeable," he suggests, because

[2] See *Studies in Ethnomethodology,* for example.

[3] *Constitution,* p25

[4] Social theorists have noted deep equivalences between Strauss's "universes of discourse", Chicago school sociology, and Giddens' "structuration, so I take myself to be using a complementary set of frameworks.

they "know a great deal about the conditions and consequences of what they do in their day-to-day lives. Such knowledge is not wholly propositional in character, nor is it incidental to their activities."[5] This fact encourages us to take the testimony of actors seriously, and leaves open the amusing possibility that academic sociological findings might turn out to be common knowledge among the subjects being studied.

On the structure side, Giddens introduces a two main ontological categories: 1) "rules" and 2) "resources." Rules provide the agent with principles of communication and interpretation and sources of "legitimation" (i.e. norms and values). Resources can be "allocative" in terms of material and economic goods or "authoritative" in terms of control over other humans. Structure does not always act alone, of course. If it did, then we might see ourselves as hopelessly determined by our social environment. But that is not the case in Giddens' framework. The agent draws on "memory traces"—STS scholars would likely prefer the language of "imaginary"—of past structure to recreate and reshape it in the present. In this way structure is not merely material or law-like patterns; it is rooted in the practical consciousness and created by purposive human action. To illustrate: if I find myself with the role of university professor, I can mimic (or even resist) what I remember from the behavior of my professors in college. I might email my students to let them know my "office hours," even though no one asked or required me to do so. I walk in the department office expecting dry-erase markers to be available for my use. The remembered rules and resources of the professor shape my actions in that role.

This idea of structuration may strike some readers as overly schematic. Smith (2005) calls it a return to "blob-ontology," a picture of sociality filled with abstract concepts that have no determinate referents. And it would be an empty explanation indeed to say blithely that neural engineering is the result of rules and resources in combination with knowledgable agents. Used in this way, Giddens' framework seems like a poor tool for understanding technoscience. But structuration and its accompanying social ontology function for me more as "sensitizing devices" than as the one "true" social theory.[6] It suggests what I can expect to learn from investigating some particular social environment and directs my attention as a qualitative researcher towards a range of phenomena and social things that could be relevant to my research. I know that I've not finished my analysis if I've found nothing that might be called a rule or a resource in Giddens' parlance.

[5] *Constitution,* p281

[6] See Duberley et al. (2006) for more about this sensitizing function.

Moreover, structuration implies a couple of important things for my research. First, technoscientific actors can tell me about the social world in which they inhabit. They are familiar with the rules and resources that enable and constrain their actions. Second, structuration suggests that technoscientists are not helpless victims of powerful social forces. On the contrary, it acknowledges the individuals' creativity (in the sense of both innovation and creation) in producing structure. But as applied to attributions of responsibility, it makes explicit social conditions that are left untheorized in Douglas' work. The upshot here is that the primary unit of responsibilist analysis cannot be the individual or collective agent in isolation. (Ir)responsible behavior in technoscience is produced by resources at hand, by contextual norms and values, and by the local rules of communication. Accordingly, a reasonable and contextually-sensitive understanding of technoscientific responsibility requires a familiarity with these factors. Merely imposing codes of ethics or bases of responsibility from outside the practice may have no effect. The knowledgeable agent engaged in technoscience, then, likely has a better understanding of what technoscientific responsibility could be than the external agent engaged in philosophy.

But can we trust what they say? Imaginaries and the value-ladeness of social worlds

But keep in mind: the "knowledgeable agent" is never a neutral agent. Though my interviewees are likely well-aware of the "rules" and "resources" at work in their practices, actors in neural engineering may also imbue their social self-understanding with normative content about the proper role of technoscience within society, about methods of accountability, and about their responsibilities to the public. Consider Douglas' and Fuller's complaint that Kuhn's vision of scientific practice made science immune from external critique, or indeed, my own complaints about philosophers of science (*including Douglas*) in the previous chapters. These are cases where a specific understanding of social reality encodes a particular set of values and desirable futures. Scholarly work on imaginaries and collective visions makes this point clearer and gives a much needed elaboration of Giddens' placeholder, "memory traces."

George Marcus (1995) was the first to emphasize that technoscientific actors understand their disciplinary innovation and trajectory in terms of "technoscientific imaginaries." In his introduction, Marcus explains how the scientists' visions of technological futures are often compared to science fiction, whether outlandish space operas or undersea utopia; And yet, Marcus suggests that these visions take on more modest and political forms in reality. Crucial in these less grand imaginaries are specific plans for how organizations should be structured and how labor will be divided up amongst technical experts. Sociologists will find similarities here to Ann Swidler's (1986) notion of "repertoire."[7] Gary Lee Downey's contribution to Marcus (1995), for instance, reveals how members of the ACSYNT institute understood themselves and their work as remedying a impractical disconnect between industry, academia, and government in the 1990s; aeronautical engineering could be improved, they argued, if it bridges these institutional worlds by bringing together actors from each sector. I will argue later in this chapter that the technoscientific imaginary at work in the CSNE contains similar connotations of novel institutional hybrids for the sake of better engineering and quicker translation.

Zooming out beyond narrowly instrumental technoscientific goals, the "sociotechnical imaginaries" of Jasanoff & Kim (2015) are not limited to the goal-oriented visions of engineers or scientists but situate technoscience within broader visions of social order and desirable futures. According to the authors, sociotechnical imaginaries:

> Occupy the theoretically undeveloped space between the idealistic collective imaginations identified by social and political theorists and the hybrid but politically neutered networks or assemblages with which STS scholars often describe reality. Our definition pulls together the normativity of the imagination with the materiality of networks: sociotechnical imaginaries thus are "collectively held and performed visions of desirable futures" (or of resistance against the undesirable)[8]

The idea is to enable conceptual bridges between, for example, the nationalist narrative of taxpayer-funded American dominance in space and the technical-institutional networks that make spacecraft and their operators possible. Or in the case of neural engineering, between widely shared conceptions of the normal human body and disability, on the one hand, and the immense network of actors that is assembled around neural devices. In both cases, the former is a shallow assertion without the latter, and the latter is opaque to critique without the former.

Whether sociotechnical or technoscientific, imaginaries are not just mental entities or fantasies. When embedded in human practice, imaginaries marshal normative visions of how the world (or some part of it) should be, motivating action and structuring our collective projects. And unlike

[7] Recently, philosophers of science have re-invented these ideas but in a crudely instrumental fashion, with no reference to the sociological precedent. See Leonelli and Ankeny. (2015) "Repertoires: How to Transform a Project into a Research Community"

[8] *Dreamscapes*, p19

Giddens' "memory traces," the framework laid out in Jasanoff & Kim (2015) allows us to situate the practical consciousness of technoscientists in a temporal and discursive context. Jasanoff and Kim argue that imaginaries may begin as "vanguard visions," as the plans or dreams of one person or a small group but eventually become something more. The authors call this transition a process of "embedding," by which the imaginary "must latch onto tangible things that circulate and generate economic or social value: commodities like wine or diamonds; artifacts such as defensive weaponry or GM crops..." It is at this stage that imaginaries are discoverable on 'About Us' pages on websites, in funding proposals, on front pages and in the president's speeches. The normative content and desirable futures are visible and active in setting societal agendas and codifying specific solutions.

Embedded imaginaries are useful in my project because they link the individual technoscientist and her technical toolset to an overarching motivation, a philosophically-evaluable set of ethical and political assumptions about what our problems are and how we should go about solving them. By observing the texts and discourses of neural engineering, my own qualitative work can begin to reveal the "neural imaginary" at work in the CSNE. But because imaginaries are not neutral descriptions of reality, it is crucial that we not take scientists and engineers at their word when they describe how the CSNE works and what makes it an appropriate answer to the perceived problems of society. Their understandings of themselves and their social environment are not untruthful, but the outsider, the scholar-analyst may have some critique to offer or a new view on the practice.

Interviews with an ethicist: data and methods

The data

The Neuroethics Thrust started self-consciously qualitative research in 2015; we arranged a second set of interviews, retooled to listen carefully rather than to intervene actively. We recruited interviewees through email, distributed via the CSNE mailing list, and by word of mouth. Many of the participants in the previous project volunteered to be interviewed again. Others had left the Center. In the end, we recruited fifteen participants, covering a range disciplines, including electrical engineering, biology, neuroscience, medicine, materials science, and others. These interviewees were spread across three institutions, University of Washington in Seattle, San Diego State University, and Massachusetts Institute of Technology. Working on my own, I was also contacted by some junior lab members—they were a mix of postdoctoral and graduate level—who offered to provide an interview. In total, we enrolled about half of the CSNE's senior faculty into our study, with a supplementary group of five junior researchers. Both sets of interviews were reviewed and approved by the University of Washington Human Subjects Division.[9] Sampling was not purposive, since we were content to have any CSNE member opt-in to the study. As a result, representation of diverse genders, ages, and ethnicities was partially a function their relative frequency in our pool of participants.

The 2015 interviews were semi-structured. The interviewers, armed with a digital audio recorder, met the interviewee in person or called the interviewee's lab. We used a written list of questions to guide the conversation and ensure uniformity across interviews. The guide inquired about many related topics: *i*) ethical "issues" already encountered by the interviewee, *ii*) the targeted "end-user" of the research, if any, *iii*) methods of communicating with the media, *iv*) the role of funding in research, *v*) key design parameters and constraints in the research, *vi*) integrating potential "end-users" into on-going research, *vii*) and (a "bonus question") the interviewees thoughts on "ethical issues" that had already been highlighted by the Neuroethics Thrust in previous reports.

Even so, our plans for the interview guide were overly optimistic about what could be covered in a short interview, speaking across disciplines, and a variety of things introduced variation into the interview content. Limited to sixty to ninety minutes of discussion, many interviews omitted the final section. Many interviewees didn't seem to understand some of the questions, despite my attempts at reframing. One of them voiced dissatisfaction with the questions themselves, saying they were "nonsense" or "no good" despite providing answers that were sociologically illuminating for me. Sometimes, I had to temporarily pause the interview, which changed the trajectory of our discussion; one PI stopped to converse with a panicked lab member, another to order a takeout lunch. Between these and other interruptions from the real world, the comparability of interviews is less than perfect, which I took into account in my analysis. Specifically, I avoid inferring correlations between variables and themes (e.g. seniority and concern for funding), which would not be reliable based on the data.

[9] Human subject applications #48498 and #50162, respectively.

With the audio recordings in hand, we hired a transcription service to convert all of the audio files into time-stamped transcripts, which each interviewer then checked for accuracy. My supplementary interviews were transcribed and edited myself. In combination, these edited transcripts comprise the raw data on which this chapter is based. I used NVivo to open-code each transcript, one-by-one, letting the coding tree develop iteratively throughout the process. This process amounted to two complete rounds of open-coding, before I shifted to a less linear process of comparing, merging, and relabeling codes to achieve a consistently coded dataset related to my research interests. It was in this focused coding that that I began to develop themes. I identified each theme by looking across interviews; matrices of frequent codes plotted against individual interviews showed that some phrasings and ideas were shared by many of the interviewees.

Coding and data analysis

Methodologically, my process of analysis was thoroughly constructivist. I had no foundationalist aspirations to make purely inductive or deductive inferences from the data, to generalize about some statistical population, or to derive testable hypotheses from existing theory. I reject the disjunction between empiricist foundationalism and relativist wishes. My own situated contributions and background assumptions made these epistemologies a poor fit; beyond my outward identity as CSNE ethicist, I specifically value perspectives of persons with disabilities and have a personal interest in preventing technocratic concentrations of power in society. Habitually, I look for promises, predictions, and utopian visions that are left unquestioned in technological development. These facts about myself are not inconsequential in the way I pick out themes from interview transcripts. But neither are they necessarily reasons to distrust or set aside my conclusions in this chapter.

There are well-established precedents for my style of reasoning in the qualitative methods literature. Most influential, for me, is Howard Becker's suggestion in "Whose Side Are We On?" Becker (1967). He notes that qualitative sociological research does not demand that we forget ourselves or our commitments. Instead, he suggests, aspiring social scientists should place themselves in a situation long enough to "have a perspective on it," using social theory to guide their attention to the full range of data. The overarching theme in his methodology is an openness to the world, in keeping with the Chicago school tradition. Of course, the result is still a perspective and not the god's-eye view; as we begin to see the world through the eyes of a lab director, we may lack the perspective of the graduate student or the funding administrator. This situated character of social understandings is not inconsequential.

Sociological work, Becker reminds us, is itself subject to "hierarchies of credibility" in which the powerful are able to dictate the authoritative description of reality, despite the plurality of perspectives or the views from below. And yet, he asserts, it is precisely the hierarchy of credibility that is at issue when we direct our scholarly gaze to political situations. Even in recusing ourselves from political contestation, we locate our research with respect to existing hierarchies. In my research, the relationship between technoscience and society is sufficiently politically fraught that any accusations of "bias" are more likely to originate from stakeholders (technoscientists or persons with disabilities) than from my peers in philosophy or STS. An unbiased view is strategically important for those with a direct stake in the practice.

I am trying to stress, here, that re-description of social worlds is inherently normative. The analyst's findings will inevitably problematize or align with the authoritative story on what our shared reality is like, especially our socio-political reality. Nevertheless, once we accept this inevitably political element in sociological findings, there are concrete ways to cultivate an openness to the world. To this end, my methods of data analysis are in keeping with guidelines laid out in "constructive grounded theory," as described by Kathy Charmaz, and "situational analysis," from Adele Clark. Both authors begin with the classic framework of Glaser and Strauss—data is for open-ended theory construction rather than the statistical representation of some population—but deny the positivist assumption that theory emerges from the data alone or that inference is simply a result of logical operations on the data.

Charmaz (2006) reframes grounded theory by placing it firmly in the constructivist, interpretive tradition. She keeps the basic formula of grounded theory: collect data, code and constantly seek out comparisons, test out different analytic categorizations, follow hunches, and seek out more data when needed. As in classic grounded theory, the iterative character of this method allows the analyst to place empirical pressure on possible interpretations of the data without having to deduce testable hypotheses from some prior theory. One starts with the data, rather than a literature review, and any preconceptions have to "earn" a place in the analysis through coding. This openness to the world, Charmaz adds, is not incompatible with the fact that enacting grounded theory is itself an active, social process *in the world*. So we must realize that sincerely open-ended qualitative work is possible thanks to—not in spite of—the contextual interactions between researcher and research subjects, the particular execution and justification of method, background commitments of the actors, and so on.

"Situational analysis" in Clarke (2005) adds a useful conceptual twist to constructivist grounded theory, and I found myself routinely using Clarke's framework to make sense of my final stages of coding. She adapts social theories of collective action (as exemplified by Gidden's "structuration", Strauss' "universes of discourse" and Becker's "doing things together") into the analytic-methodological unit "social world." In practical terms, this suggestion amounts to making graphical representations, maps, of the social world being studied, marking boundaries, relations, and surprising absences. Clarke specifies discourses, whether from texts or interviews, as particularly fruitful sources to draw upon. The map functions to answer a set of questions about collective action, including:

- What is the work of each world?
- What are the commitments of a given world?
- How does the world describe itself in its discourses?
- How does it describe other worlds in the arena?
- What actions have been taken in the past and are anticipated in the future?
- How is the work of furthering that social world's agenda organized?
- Are there particular sites where the action is organized?[10]

In answering these questions, mapped social worlds may point to unmapped areas, missing data, and future projects for the analyst. As a corollary, you know your map is in good shape when no new social worlds are forthcoming from the data. In reading Clarke, my themes suddenly took on a new significance, serving as content in my emerging map of neural engineering.

Charmaz and Clarke provided me sufficient methodological guidance to make sense of my interviews and a toolset appropriate for politically-charged but open-ended qualitative research. In the terms of Clarke and Star (2008), I had found a "theory/methods package." By following their

[10] see p115 of *Situational Analysis*

suggestions, my conclusions are likely what Merton would label "middle range" theory; they are not overwhelmed by the messy minutiae of the data, but are less abstract than universal or all-purpose social theory. In his words, they are "theories that lie between the minor but necessary working hypotheses that evolve in abundance during day-to-day research and the all-inclusive systematic efforts to develop a unified theory that will explain all the observed uniformities of social behavior, social organization, and social change."[11] In my case, I am satisfied with framing my results as "the social world of neural engineers," an attempt to leverage discourse and map out a subsection of present-day technoscience as understood by its practitioners.

Ambiguous disciplines, amorphous publics: qualitative findings

The methodological and theoretical literature introduced above suggests that I can learn from my interviewees in a way that doesn't set normativity aside. I hope now, after that extensive preamble, the reader can understand this assertion: the qualitative work presented below introduces a sociotechnical imaginary that is "embedded" in neural engineering. The imaginary reveals for the outsider the social roles, external social worlds, desired futures, and internal contradictions that define neural engineering practice and guide actors as they navigate their social environment. More concretely, I will describe how neural engineers see themselves as negotiating with industry and the state, creating new institutional formations, and ultimately bringing about a neural future where inequality is remedied through biomedical interventions.

As I hope to show, this neural imaginary explains how my interviewees could all perform the same optimistic vision of a technologically-enabled egalitarian future while expressing profound doubts about its socio-technical feasibility. Equally significant, the dominant sociotechnical imaginary reveals the way in which neural engineering effectively (though perhaps unintentionally) disenfranchises the publics that it is meant to benefit.

Sketching a social world

Following Clarke's methodological prescription, I sketched out the neural engineering arena from the perspective of CSNE researchers (Figure 3). Scale, here, is not quantitatively or even qualita-

[11] Merton (1968)

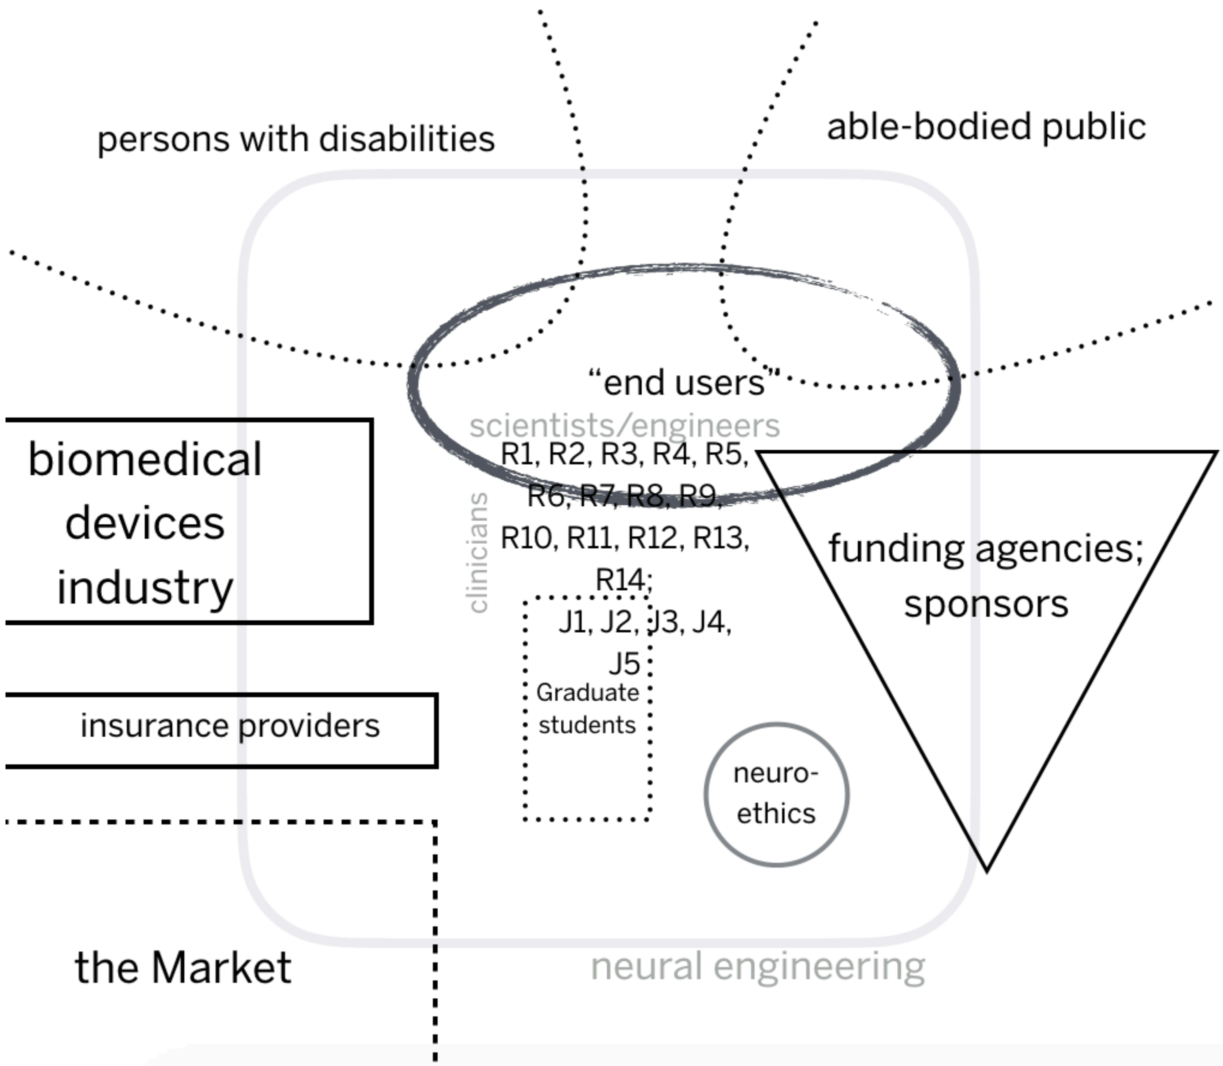

Figure 3: The Social World of Neural Engineering

tively precise, but the presence and position of actors and overlapping social worlds is significant. Immediately, the diagram pushes us to reconceptualize what counts as neural engineering, what is inside and what is outside. The reader may notice, however, that discourses themselves are not part of the map; they are absent because the analysts' description of social worlds are, in a sense, a distillation of discourse. Social worlds supervene, we might say, on discourse. Thus, what is being illustrated here is not "the" world but the world that is meaningful to my interview participants. The map would look drastically different from the point of view of persons with disabilities or from the philosopher's armchair.

The arena revealed by my data implicates a range of social worlds. The medical devices industry is present, along with the state and its funding agencies like DARPA, NSF, and NIH. The social world of neural engineering itself sits somewhere in between those of science, engineering, medicine, and university pedagogy. Perhaps less obvious is the prominence of graduate students or technoscientists-in-training, and the presence of neuroethics. The "end user" or public is plural, encompassing both the able-bodied "consumer" and persons with disability, as well as funding agencies or other researchers. Lastly, "the market" and insurance companies often appear when discussing questions of access to technology.

The arena laid out here represents one of the basic functions of the sociotechnical imaginary; the imaginary gives members of a collective a sense of how to plan and justify their day-to-day actions. We see here some of the social worlds that neural engineers see as external yet nonetheless present. Neuroscientists may not, for example, desire interaction with a federal agency but negotiation may be a matter of necessity. Overlapping publics make an appearance though for different reasons. The map alone, however, is not enough to understand the neural imaginary present in my interview data. By diving into the more detailed relations and properties of this map, we see also accountabilities, collaborations, desired futures, and underlying doubts. The next few sections will begin to do exactly this.

The scientist-engineer and the limits of interdisciplinarity

There is a cautionary tale lurking in the way "transdisciplinarity" (from Gibbons et al. (1994)) evolved from an analyst's category to a best practice in the governance of science and technology. Regardless, my own data was filled with images of merging disciplines, interdisciplinary collaborations, and boundary traversal. One subset of researchers expressed that realizing the potential of neural technology was only possible with new institutional spaces, where research was both "basic and applied," discovery-oriented and clinically-relevant:

> I think that the real strength of the center is the interdisciplinary approach bringing all these people from different fields together and I think that it's clear that, you know, that kind of effort is what it's going to take in the field of neural engineering to actually come up with a clinical device that's useful. The fact that we have all those people in one center, working on that, I think is another asset that gives me optimism that we'll actually be able to make impact.

As within the ACSYNT institute described in Marcus (1995), members see their own work as useful because it is integrated in the right sort of institutional structure, a coalition of engineers, scientists, clinicians and industry representatives, all funded by the National Science Foundation. The very meaning of neural engineering seems to be tied to this vision; almost all interviewees talked about neural engineering as a "mixture," "combination," or "everything from fundamental engineering to fundamental science." Sometimes this characterization came to define researchers' identity: one of the interviewees claimed to be neither an engineer nor a scientist but rather an "ambassador" between the two. Meanwhile, I could find no explicit attempts to preserve "basic science" as a practice distinct from others or separated from zones of application.

These findings should not be too surprising given parallel language in the NSF solicitations for ERC funding. As I mentioned in the previous chapter, research at the CSNE is described by the NSF as "systems-motivated basic research." I suspect that broadening the scope of my qualitative research would find traces of this epistemology across funding agencies, journals, and international research clusters. There is an philosophical question embedded here regarding the life of knowledge claims made in such spaces. How can an engineer's practical needs dovetail with the evidentiary standards of cell biology? What is the balance between opening projects to peer review and closing them off in order to maintain intellectual property rights?[12] Whatever the answer, the implicit epistemology of places like the CSNE provides a de facto definition of knowledge; knowledge is something that can successfully travel between molecular biology labs, clinics, micro-fabrication facilities, and corporate boardroom meetings. At the end of the day, it is worth asking if these particular translations and "trading zones" are appropriate, whether for fulfilling the vision of a neural-technological utopia or for realizing our more abstract ideals of democratic, critical knowledge practices.

Interestingly, while many CSNE researchers were happy to traverse disciplinary boundaries between science and engineering, some researchers were clear that they were "not an ethicist" or that they relied on the neuroethics thrust to work on potential socio-ethical problems of CSNE research. The reasons behind this reluctance to cross into the arena of neuroethics varied and were not alway clear to me. Sometimes, it was a matter of "job description: "It [ethics work] doesn't seem like the job description you actually work on. It doesn't seem part of the job. It's

[12]It is worth noting that in the course of my research I've personally been left off of email chains and politely asked to skip meetings so that I am not exposed to patentable ideas.

not the message that you get. Does that make sense?" This interviewee said that they would "outsource" ethical reasoning just as if they had a "really hard computer problem," which would involve a mathematician. In other responses, competence or expertise was cited as a determining factor for who considers socio-ethical questions: "we, within our lab are not... we don't have the expertise or the experience even probably know what the right questions are, so we have to go out." Resources, too, were cited as too scarce to allow for ethics on the clock or in the lab.

A remarkably frequent association with this theme of "not doing ethics" was the claim that the interviewee's research is at too early of a stage, too narrowly focused, to think about the ethics or societal impact of neural devices: "To be honest, I think the biggest ethical issues are going to not be in my hands, they're going to be in the hands of the clinicians that are using these devices for people, and selecting the people that will actually use them." And, "There are also ethical issues in terms of a long-term translational potential for this work and I suspect you'll wanna hear more about that, but we are pretty far away from a clinical device. [...] We don't think about those particular issues very deeply at this point. We know that they're there." This statement really puzzled me, since discourse at the CSNE was so shot through with visions of how neural devices might work and how they could bring so many benefits of a certain type; ethical reasoning or moral imagination was already at work! But my conversations at the Center led me to believe that this vision was somehow frozen, set aside and unquestionable until *n* amount of basic technical work was done.

It is notable that even though the socio-ethical issues we discussed frequently implicated persons with disabilities or members of the general public, their concerns were repeatedly taken to be within the purview of ethicists or clinicians. The CSNE is "designed" to incorporate ethical reasoning, since these experts are present. I found this to be striking considering that our interview guide specifically asked how to integrate "end user input" into the research and design process, a wording which implies more direct forms of contact between end users and researchers. Yet, it is to the advantage of actors in technoscience to isolate themselves from the public via proxies of various sorts. The fact that ethical "experts" can play this role has already been described by Evans (2006) in reference to institutionalized bioethics and by Doubleday et al (2010) in reference to embedded social scientists. In both pieces, the authors stress the difficulty of personally fulfilling that role and push us to question who is allowed to speak on behalf of the intended beneficiaries or stakeholders of technoscientific research. Philosophers of science should pay attention to the equivocations and ambiguities around "the public" and its representation within technoscience.

Distributed labor and unpredictable outcomes

The idea of collaboration, interdependence, and "team effort" dominated researchers understanding of neural engineering. Much of this language was associated with our questions about who might be responsible for technology success or failure and whether or not the researchers were in control of their projects. The most common response on this topic was that the CSNE's work is a "team effort," such that no one person is in full control and that some responsibilities are shared: "We all bear the same responsibility. I think it's a shared... either way, it's a shared experience and that's what collaboration is all about," and "It's a huge team effort." Even when we inquired whether more individual responsibility would be desirable, some interviewees saw the "team" structure as unavoidable:

> Interviewer: Do you think that engineers and scientists should strive to create conditions where researchers have more individual freedom to conduct their research and translation?
> Interviewee: Not necessarily, the reason being when you have a team of logical people working, just about any transformational research will need teams of diverse people, that's unavoidable. When you're working with a team of diverse people, it is a team effort and just those people are logical and they're all are coming towards [the] common unified goal.

As you might expect, this idea of teamwork gave some researchers a sense of purpose in their individual work. One researcher used the language of "service:"

> I provide the research more like a service to, you know, to other researchers. So my research really is intended to, you know, enhance and improve and further the research of, you know, [interviewee colleague at CSNE] or [interviewee colleague at CSNE] or, you know, [interviewee colleague at CSNE], so the people that are working with patients that are, you know, advancing. And I think I'm not independent.

In this case, the service functions to advance research of colleagues that work closer to patients (e.g. clinicians or surgeons). This was a refrain—research is channeled through clinical experts—that I encountered frequently.

The infrastructural contents of the neural imaginary are not limited to the "team effort," however. Interviewees frequently drew my attention to a vast network extending well outside of the walls of the CSNE. Some of these instances were in response to our question about "broader social forces" that shape research, but others occurred more spontaneously in conversation. Funding and its sources were, of course, a dominant concern. The majority of researchers presented their research projects and their output as a factor of what is fundable, with a subset of complaints about how funding follows "trends" and doesn't correspond to actual need. And though researchers individually "bought into" the NSF-CSNE plan to create principles of brain-computer interfaces, National Science Foundation administrators have some control over how the plan is implemented via yearly "site visits." From my own experience attending them, the annual visits came with a collective sense of existential crisis; if the NSF was not impressed with the CSNE's progress, the Center could lose its funding.

Abstract invocations of "the market," "industry," and "insurance" also made several appearances in interviews, treated as if they had an agency of their own. One researcher was already strategizing how to get around market obstacles:

> My biggest concerns are that the populations we're trying to assist are by nature small which is a great thing; it's a great problem to have, but it means economically there's not always a market to support the complexity of the devices that we're trying to design. And so I'm always looking to broaden that market from people with spinal injury, to people with stroke, with traumatic brain injury or beyond.

As illustrated by this response, some researchers see the market as controlling access to what devices get made and who ultimately receives them. Concerns around healthcare insurance and the medical devices industry followed a similar formula; they can affect how the technology is taken up (or not) in a way that researchers cannot control.

When discussing the distributed character of neural tech development, interviewees often stressed the ways in which technological development was unpredictable or out of their individual purview. Between the market, industry collaborators, insurance companies, federal regulatory agencies, funding trends, among other things, I was given a picture of neural engineering as a cacophony of overlapping social worlds, some of which do not share the goal of creating brain-computer interfaces, per se. One interviewee cited a previous personal experience where a project was saved from cancellation only by the passionate interventions of patient-activists who wanted to overcome obstacles to commercialization. I could not help but be sympathetic in hearing these stories, as interviewees tried to help us understand the difficulties of navigating in the neural arena. Yet, I also worried that sharing the researchers' imaginary or seeing neural engineering as distributed out into society would eventually rob me of my ability to critically evaluate neural engineering or to hold anyone accountable.

Researchers' pessimism became most intense, as you might expect, when we discussed the possibility of failure. We asked the interviewees if they were responsible for a failed device translation. Sometimes they claimed some individual responsibility, letting a somber silence follow. But a substantial majority were filled with new energy at the end of the interview to reconceptualize failure. A couple of interviewees said that failure-to-translate is to be expected, since "that that's like 90% of medical innovation" and "we know that most medical devices don't make it to translation, we know that startups don't survive, I mean there's many valleys of death along the way, so the reality is it's unlikely that given this, that we'll make it all way through." But even in these cases where a device is not produced, the researchers are keen to point out that failure to translate isn't really failure:

> You know, we never do research basically with the intent that it will eventually be successful, right? Whatever we do, you know, will eventually– you know, will not disappear, will transfer it into some other research. It will feed into other research and will enable other research. [E]ven if they don't launch basically at this stage into marketable products and, you know, from our side we'll still feel that, you know, the outcomes that we obtain will translate into other research and eventually maybe not today's test beds but to the next generation test beds will basically be successful.

After all, they reminded us, learning how *not* to make or translate a technology is itself a type of success.

Philosophers can gain new insights from appreciating this distributed vision of technoscience. Let's take it to be roughly adequate to say that technoscience is not separated from society but rather shades into it through myriad institutions and organizations. To the extent that all technoscience is indeed as its practitioners say, it showcases how epistemic and technical success in technoscience works; far from wishing-making-true, success requires an immense assemblage of materials, actors, and institutions. A rich understanding of social construction does not imply that true things are less true or that there is no world pushing back against our technoscientific endeavors. This feature suggests new types of normative epistemologies for technoscience. But to return to a more critical angle, the neural imaginary can also determine our ideas of responsibility and accountability in neural engineering. Consider, for example, that we might not assign much responsibility to neural engineers if they are really just one node in a network. Philosophers and STS scholars have an opportunity to inquire into the implications of this vision and if it is mirrored in more comprehensive qualitative work.

Fixing Disability and Uncertain Publics

We asked participants about the "end users" of research, wondering who they might be in the minds of researchers. In my analysis, I found that most participants express a strong commitment to meeting the needs of people with spinal cord injuries, stroke, and related neurological conditions. These target populations, mostly persons with some form of disability, were frequently described as people who "need", "want", or don't know that they need neural technology. One researcher provided the strikingly categorical statement: "[Patients] wish to have a new technology; they want to improve their lives so they want to take the new technology." Other interviews were more descriptive about why that might be the case. Another researcher conducted a sort of thought experiment and settled on independence: "I think independence is key. So, having control over your own body, being able to do things on your own at least to me would feel like that would be quite enabling." Many others settled on reasons centered on medical notions of recovering some neurological "function" and "functionality:" "Think how much better the life of a person can be when he cannot move the limbs, will be if we can just capture his waves, his thoughts, and make... and give him access to his upper limb, that's the focus right now, the upper limbs." A medical and individualist model of disability corresponds to the subject matter of neural engineering is often the body or, more specifically, neural pathways that have become disconnected.

What is surprising, perhaps, is how insistent some researchers were about the potential to improve lives with neural technology despite noting misgivings among persons with disabilities or individuals who have "adapted" to their injuries. We heard the admission: "I'm not sure the end user knows...not necessarily...what's best for them." The interviewee proceeded to explain that if someone has already adapted to a condition, like paralysis, they won't see the necessity of invasive neural technology, which might seem attractive to someone who was just injured. This thread cropped up again with a different PI:

> It seems like people who have had their injury [...] or impairments longer, I think have adjusted more completely to that. And I think also may have seen some of the trends come and go, and seen things that look very helpful, will not come to fruition. And probably have quite frankly a bit more realistic, if not pessimistic, but certainly more realistic picture of how this technology could fail to impact people. And I think people who have had these impairments for shorter periods of time have not been tuned in to the winds of change, I think I know that as much, and are probably at different stages in adjustment to disability as well."

Here, the PI also seems to worry that some potential end users have become unreasonably cynical in the face of ineffective biomedical technologies. Even this more sober assessment was accompanied by a confident assertion that the CSNE would eventually lead to benefit, "absolutely."

Unpacking these statements could go in several productive directions. We could ask what sort of ideal of the human body is presumed—surely there is some valuation here about being an ambulatory non-paralyzed person—but that doesn't fully capture the notions of "helping" and "improving." "The center's goal is to develop ways to help people with neurological conditions such as stroke and spinal cord injury," says Director Raj Rao in his mission statement. Implicit in these statements is not only an ideal of the human body, but also a biological model of disability and an egalitarianism that requires differential treatment for the differently able. Disability is to be erased, bodies would be made equal, not by fixing the social structures that preference one sort of body over another but by intervening on the bodies of persons with disabilities. This commitment to fixing disability seems to be a rule of legitimation within neural engineering, almost to the exclusion of other possible commitments. Relatively scarce in many of the interviews was the idea of enhancement, of modifying anyones body to augment function. One PI suggested that enhancement of the able-bodied public is not something that they can easily talk about, since "Congress is not gonna like it and NIH is not gonna like it. NSF is not gonna like it."

The content of neural engineering's egalitarianism of bodies becomes more explicit if we compare it briefly to the way that many activists and disability studies scholars portray disability. A full investigation of counter-imaginaries is beyond the scope of this project, but a few examples suggest the richness of such a comparison. Since at least the 1970s, many disability activists have challenged the idea that biological or bodily difference is the relevant cause of disability. The Union of the Physically Impaired against Segregation (UPIAS) in the UK asserted as a fundamental principle, "It is society which disables physically impaired people. Disability is something imposed on top of our impairments, by the way we are unnecessarily isolated and excluded from full participation in society. Disabled people are therefore an oppressed group in society."[13] More recently, Sue Austin's (2012) performance piece, *Creating the Spectacle*, challenges the viewer with images of a woman scuba diving in a wheelchair, floating serenely in a neutrally-buoyant state; far from being bound by her chair, she seems freer than most able-bodied people walking on land. Disability in these depictions is not about the body but about the situation. Consequently, if we want to reduce the disadvantages associated with being disabled, then our attention should be directed towards social situations and not towards the fact of someone's using a wheelchair or not being able to voluntarily control hand movements.

Anita Silvers (1994) shows the implications of ignoring this, the social model of disability, in her criticism of Amartya Sen and like-minded egalitarians. There is a fatal flaw, she suggests, in equivocating between "being normal" and "being equal." Able-bodied citizens have difficulty imagining the value in life without walking upright or without four functional limbs, so it becomes almost impossible to give persons with atypical bodies a place in our evolving conceptions of egalitarian society. Silvers points out the resulting anomaly in Sen's work by which he awards persons with disability almost unlimited resources to repair themselves to be "normal" in the name of equality, but then limited resources for everything else. This result is not an idiosyncrasy of the political realm, unfortunately. Goering (2008) argues that academic (bio)ethical judgments of value are just as clouded by this form of equivocation. The effect is virtually the same, persons with disability are not even given the chance to deny compensation by society or to assert the value of their form of life. They don't even have sufficient standing to turn down "help."

Had I thought about this apparent moral/political mismatch prior to my interviews, perhaps I would have not been surprised to learn that most researchers do not have frequent contact with potential end users. Though a majority of researchers agreed that "end user input" was valuable at some stage in research and development, only a small subset claimed to interact with potential "end users" on a regular basis. When I asked about their contact with people who might be part of the CSNE's target population, most of my interviewees cited unplanned or indirect exposure:

[13] 1976, p4

"Not directly. I have, you know– not directly with anybody, but sometimes like I watched, you know, demonstrations and like I read interviews, but I did not interact with anyone that might potentially be using, you know, use this implantable technology." A few researchers hesitated to generalize about end user needs. These statements are somewhat at odds with the stated aim of the CSNE to benefit persons with disabilities. One might expect that the people claiming to meet the needs of persons with disabilities would be more fluent in the lived experiences that might necessitate neural devices. But as long as persons with the targeted disability are present as vague ideas (rather than as persons), it is difficult to see how they are fairly represented in neural engineering.

The neural imaginary verges on incoherence when researchers redefine the intended beneficiary, re-classifying the "end-user" to also include other researchers, the able-bodied public, and even an unknown placeholder. As mentioned in the above excerpt about research as "service," some interviewees see their colleagues in neural engineers as "end users" or "intermediate barriers" to the patient. Some CSNE members are "closer" to the patient in virtue of being a neurosurgeon or clinician who might one day implant the eventual neural device, while others are the stereotypical "scientists" who would pursue any new findings out of curiousity. A few interviewees also admitted that the "end users" may turn out to be able-bodied: "BCIs might provide a real – not a temporary fix, but this could be a real solution for people with stroke. But I like to think that we also think about everyone else and not as a medically assistive device, but more as an assistive consumer technology like cellphones and computers." Though provocative, these rare statements did not give me a clear idea of what BCIs could do for the able-bodied. The ultimate aim of the CSNE gets even fuzzier when researchers say things like, "'We don't know' is the short answer to who our end users are gonna be," and "the impact is gonna be distributed across multiple populations, not necessarily just the spinal cord injury and stroke."

The interviews thus present a volatile combination of desirable futures, institutional solutions, and pessimism. On the one hand, the dominant neural imaginary ensures that CSNE members express a strong consensus on the intended effect of BCIs and the underlying egalitarianism of bodies. On the other hand, there is an undercurrent of doubt about what neural tech ultimately will do in society and who it will benefit. At the same time, the very distinction between the "end users" and creators of neural devices is dissolved in some interviewee responses. I notice here a troubling similarity to the ambiguous "public" represented by our very own neuroethics team at the CSNE. I suspect there is more than a mere coincidence in these tensions; the question of publics (actual or claimed) is a core fault-line revealed by this and other qualitative work on technoscience.

Conclusion: a neural imaginary and its normative/philosophical dimensions

Since Gieryn (1983) first popularized the concept of "boundary work," the idea has served to explain the many ways in which scientists differentiate themselves from society, isolating their practice. It is a continual negotiation that has implicated even institutional bioethics, in order to keep unwieldy publics out. Kelly (2003), for example, explains how public bioethics bodies function as "border guards" to narrow the range of ethical concerns about emerging science while also limiting who is counted as a legitimate public (and who is excluded as an interest group). The effect of these bodies is to distance biomedical research from the heterogenous publics, thus maintaining the autonomy and integrity of scientific practice. Even Merton's famous norms of CUDOS can be read as functioning to distinguish scientists from your average citizen (who may not be be disinterested and communally-minded) and creating a separate practice worthy of the name Science. My own contribution to this literature bucks the trend, so to speak. My technoscientific interviewees portray themselves so thoroughly embedded *in* society (i.e. beholden to funding agencies, subject to markets, etc) that they cannot possibly operate by internal norms alone. Neural engineering, as they describe it, almost disappears the individual scientist-engineer in a distributed network of research funding, translation, and commercialization. They often see one another as "end users" and despair at the realities of capitalist markets. Regardless of whether we should share this imaginary as expressed, the end result is the same as in the old Kuhnian caricature of practice; we should find it increasingly difficult to pin any responsibility for neural technology on researchers themselves.

At the same time, my findings also reveal a tension in the self-understandings of neural engineers. On the one hand, there is a persistent optimism about helping persons with disabilities, about making bodies whole again. On the other, there is a pessimistic thread present in the interviews. Even among PIs who claim optimism, my attention was directed to unfair markets, to a capricious arational funding climate, and to the absence of persons with disabilities in the tech development process. Is this just a case of multiple personalities in each researcher? Is hypocrisy the only way to make sense of these contradictions?

I suggest we see this tension as the action of a sociotechnical imaginary, a neural imaginary, that is extensively deployed in United States research contexts and that organizes collective practices like the CSNE. In terms of Jasanoff and Kim, the imaginary is at the "embedded" stage. It is not reinvented by each person who decides to become a neural engineer; the imaginary is sufficiently dispersed in calls for funding and in federal administrative bodies that it does not rely on the "vanguard vision" of any one person (though it may have started out as such). Individuals likely encounter the imaginary as they look for funding, interact with their peers, or speak with members of tech media. It is at this embedded stage that we can look for signs of resistance or doubt with respect to the imaginary; I interpret my interviewee's pessimistic admissions in this way.

By sketching out the technoscientific and sociotechnical imaginary that sustains neural engineering practice, this chapter contributes a rather detailed story about one corner of technoscientific practice. I see this sort of work as useful on its own; not every story has to have an overt utility, and the act of re-describing enables new reflection on neural engineering. Even so, this empirical work brings important normative-philosophical dimensions to the fore. For each component of the neural imaginary, there are accompanying philosophical issues. Consider: Neural engineering at the CSNE comes with assumptions about what sort of institutional structures are needed and what disciplinary relations are appropriate to create useful technologies. Social epistemologists should find themselves questioning what sort of knowledge claims come out of these "transdisciplinary" spaces and what ethical/political values they represent (or fail to represent). Or similarly: CSNE researchers conceive of neural engineering as distributed, such that responsibility is shared. What implications does this claim (warranted or not) have for the accountability of technoscientific actors? Lastly, neural engineering is conducted on the pretense of helping a public (i.e. persons with disability), but relies on a form of egalitarianism that may not be shared by its intended beneficiaries. These are only some of the new normative fault lines at which philosophers of science can work in the co-productionist mode.

References

Becker, H. S. (1967). Whose side are we on? *Social problems*, *14*(3), 239–247.

Charmaz, K. (2006). *Constructing Grounded Theory: A Practical Guide Through Qualitative Analysis*. London: SAGE.

Clarke, A. (2005). *Situational analysis: Grounded Theory after the Postmodern Turn*. London: SAGE.

Clarke, A. E. & Star, S. L. (2008). The social worlds framework: A theory/methods package. In E. J. Hackett, O. Amsterdamska, M. Lynch, & J. Wajcman (Eds.), *The Handbook of Science & Technology Studies*, volume 3 (pp. 113–137). Cambridge: MIT Press.

Duberley, J., Cohen, L., & Mallon, M. (2006). Constructing scientific careers: Change, continuity and context. *Organization Studies*, *27*(8), 1131–1151.

Evans, J. H. (2006). Between technocracy and democratic legitimation: a proposed compromise position for common morality public bioethics. *Journal of Medicine and Philosophy*, *31*(3), 213–234.

Gibbons, M., Limoges, C., Nowotny, H., Schwartzman, S., Scott, P., & Trow, M. (1994). *The New Production of Knowledge: The Dynamics of Science and Research in Contemporary Societies*. London: Sage.

Giddens, A. (1984). *The Constitution of Society: Outline of the Theory of Structuration*. Cambridge: Polity.

Gieryn, T. F. (1983). Boundary-work and the demarcation of science from non-science: Strains and interests in professional ideologies of scientists. *American Sociological Review*, *48*(6), 781–795.

Goering, S. (2008). 'You say you're happy, but...': Contested quality of life judgments in bioethics and disability studies. *Journal of Bioethical Inquiry*, *5*(2-3), 125–135.

Jasanoff, S. & Kim, S.-H. (2015). *Dreamscapes of Modernity: Sociotechnical Imaginaries and the Fabrication of Power*. Chicago: University of Chicago Press.

Kelly, S. E. (2003). Public Bioethics and Publics: Consensus, Boundaries, and Participation in Biomedical Science Policy. *Science, Technology & Human Values*, *28*(3), 339–364.

Kjolberg, K & Wickson, F E (2010). *Nano Meets Macro: Social Perspectives on Nanoscale Sciences and Technologies*. Pan Stanford.

Leonelli, S. & Ankeny, R. A. (2015). Repertoires: how to transform a project into a research community. *BioScience*, *65*(7), 701–708.

Marcus, G. E. (1995). *Technoscientific Imaginaries: Conversations, Profiles, and Memoirs*, volume 2. Chicago: University of Chicago Press.

Merton, R. K. (1968). *Social Theory and Social Structure*. New York: Simon and Schuster.

Smith, D. E. (1987). *The Everyday World as Problematic: A Feminist Sociology*. Toronto: University of Toronto Press.

——— (2005). *Institutional Ethnography: A Sociology for People*. Lanham: Altamira.

Swidler, A. (1986). Culture in action: Symbols and strategies. *American sociological review*, *51*(2), 273–286.

Chapter 5 - Situating Good Technoscience

Across the previous few chapters, my aim has been to make the case for explicitly co-productionist philosophy of science by pointing out bad philosophical habits and misleading assumptions about how we philosophers should do our work. Simply reading some of my qualitative findings or, for that matter, any reasonably "thick"[1] description of technoscience should bring into relief what's absent in many philosophical works: concrete things like the NSF, the Bayh-Dole Act, as well as social-theoretical things, like identities, imaginaries, and constitutions. This contrast, in combination with my more detailed treatments of philosophers like Heather Douglas and Helen Longino, amounts to a series of piecemeal critiques, but in order to make a more coherent philosophical recommendation, it's necessary to collect my concerns into a couple of broad categories.

The first set of bad philosophical habits are related to divisions of labor; in what seems to be a poor mimicry of the natural sciences, philosophers specialize ever farther into epistemology, ethics, or political philosophy. I highlighted this tendency in Longino's epistemology. In reality, in practice, the various dimensions of normativity are not easily separable. Truth or Objectivity, as lone ideals for epistemology, fail to capture many problematic cases arising from technoscience. Participatory gaps often have a political valence, as when an engineer might wrongly assume a person who uses a wheelchair wants to walk again above all else, or when disability activists ask for "nothing about us, without us."[2] Scientific reasoning, even when highly logical or formal, takes on an ethical character when we think about the consequences of being wrong; that was the point of the inductive risk literature.

To dismiss these observations as social constructivism or as "not philosophy" is to ignore the actual content of the co-productionist idiom; the ways we know about the world are dependent on the ways we wish to live in the world, to paraphrase *States of Knowledge*. Neglecting this fact leads to odd pairs of papers about, for instance, why democracy demands participatory knowledge practices and how epistemology prescribes democratic social structures. The choice about which dimension of normativity is fundamental often depends on what each philosopher puts on their CV as area of specialization. This sort of dilemma should not be solved by figuring out "who is right" (e.g. the epistemologist or the political theorist). That response would return us to futile

[1] in Geertz' terms

[2] The slogan has hazy origins, but see Charlton (1998) for one account.

foundationalist debate. Instead, we can take these mirrored philosophical projects as reasons to take co-production seriously, cultivating an awareness of how scholarship too co-produces knowledge and social order.

The second set of bad habits has to do with an unreflective reliance on idiosyncratic imaginaries. We are tempted to use our disciplinary intuition as philosophers to guess what science or techno-science is like and theorize from there. Ideal theory has a long tradition in philosophy, but the habit becomes less innocent when it functions to shield scientists from criticism or prevents our attributions of responsibility from having practical relevance. I argued that Code and Douglas, with their responsibilist approach, fell into this trap for lack of sufficient engagement with techno-scientific practice. They used their limited pages and energy to address internal, disciplinary worries of philosophy rather than engaging with close, qualitative descriptions of technoscience. My work in Chapters 2 through 4 is an attempt to show through comparison how impoverished our imagination becomes if we rely unthinkingly on the sociotechnical imaginaries of Polanyi or Kuhn instead of looking at what actually sustains practice. We have a choice in how idealize, and many philosophers of science may not realize this until they leave the department or reach out to neighboring fields.

These worries could be developed further, with no end in sight; philosophy of science journals are filled with toy puzzles and reductive perspectives, some borrowed directly from the sciences themselves. One can begin to see the scope of this problem by considering the many and diverse movements that have attempted to critique the self-limiting tendencies of professional philosophy from within; these include but are not limited to history of philosophy of science (HOPOS), feminist philosophy, critical race theory, and pragmatism. This chapter, however, is dedicated to moving beyond critique and makes a positive contribution to our understanding of "good techno-science," a normative account that fulfills philosophy's prescriptive aspirations but in a way that acknowledges the idiom of co-production. My suggestion for philosophers of science is, in short, to begin with actual practices rather than conceptual abstractions or old logical paradoxes. After all, it is in the lived experience of social beings that questions of knowledge and value are collectively worked out, and neighboring disciplines provide us ample theories and methods with which to study this process. So leave the philosophy department and head for the field! Normativity will follow. Of course, the devil is in the details.

Elevator words and pragmatist alternatives

At the outset of this project, I called for a theory of "good technoscience," a corrective to the bad habits I see in other normative accounts. The vaguely positive label, "good," is intentional. To evaluate a practice in terms of "goodness" avoids, I hope, any obvious preference for one of the "elevator words,"[3] like Truth or Justice. Consciously directing our philosophical inquiry towards good or not-so-good technoscience reflects a methodological openness to various normative dimensions and should correspond to a broadly critical disposition. Equally crucial, "good" will be a thoroughly contextual matter, but as I intend it, there is no space here for the 'merely descriptive', radically relativist, or stereotypically postmodern. In this sense, my proposal bears resemblance to the American pragmatist tradition and some philosophical naturalists, like "left wing" Sellarsians. Let's rehearse a bit of John Dewey to make this point clearer.

I take Dewey's "reconstruction" of philosophy quite seriously. The diagnosis? "Hypostatization," of the good, the beautiful, and the true, into "ontological absolutes."[4] Philosophers mistake the abstract ideals driving human action for entities that exist independently from any specific context of action or project. Hypostatization of this type is stabilized by myriad institutions and materials. Philosophers write books on "epistemology of science" and hiring committees call for an applicant "who does ethics," as if such boundaries are given in advance. There is a solution to this problem, at least intellectually and methodologically; re-situate value and matters of fact in the practices in which they are active. Across Dewey's treatment of epistemology and ethics, you can see a consistent commitment to contextualism. In *Theory of Valuation* (1922), he acknowledges that some "good" things are found in common across various human operations, while stressing such patterns cannot be taken for universal principles. Resisting the urge to hypostatize forces us to pay attention to the messy details of human action, technoscience included.

"Ok, so the good is relative to context, but by what criterion do we use to evaluate within a context?" a confused philosopher might ask. "It seems you are confusing philosophy's project with description." Dewey's proposal resists specifying what the good is in terms of a criterion and instead emphasizes the on-going process—he calls it "experimental"—by which we examine our current "problematic situations," gather facts, and envision various actions that could reconfig-

[3]as Ian Hacking would say
[4]*Theory of Inquiry* (1938), p177

ure our original situation. The pragmatic method,

> ...says that it is good to reflect upon an act in terms of its consequences and act upon the reflection. For the consequences disclosed will make possible a better judgment of good. Thus the good of foreseen consequences is not final nor dogmatically determined. It is good as a 'better than'—better than would exist if judgment had not intervened.[5]

It is the open character of the good, the fact that the good changes in light of new discoveries and insights, that gives "experimental inquiry" its essentially normative dimension and its capacity to replace traditional theories of value. As long as we humans see ourselves as beings with needs and projects, the pragmatic conception of normativity is better way to explicate "the good" in its many contextual instantiations.

This picture of the good has implications for doing philosophy of technoscience. When the lone philosopher wants to evaluate some technoscientific practice, she must place herself in the context, interviewing stakeholders, mapping out their values and desirable futures, as well as observing material or institutional features of the situation. In other words, we need to talk to people, participate in their day-to-day lives, and get to know the social worlds in which they travel. This mode of research allows the philosopher to take part in the moral imagination that drives the technoscientific practice in question, which may be accessed via the lab, the Supreme Court, technology blogs, and a variety of other spaces. Knowing where to do experimental philosophy is less important than knowing how. The mode of normative evaluation—some philosophers will object to this as fanciful—is form of "dramatic rehearsal," a socially performed re-considering or re-opening of the manifold of hypothetical solutions or institutional arrangements that may meet our prior needs.[6] There is no therapeutic injection of Kant or any other moral or political theory. It is not a matter of sprinkling "public values" throughout. The good is created collaboratively in the process, and we may discover new integrative values or hidden contradictions in vision.

Dewey's reconstructive prescriptions imply that the philosopher can and should engage in qualitative research with an eye towards explicating the good. We need not pretend that ethnographies and interviews have no effect on the people we study or that our re-descriptions of technoscience

[5] "Theory of Valuation and Experimental Knowledge", p351

[6] see Alexander (1993) for more on rehearsal. Alexander uses the example of two children fighting over who is entitled to play with a baseball. We may be able to apply some presupposed principle, like property ownership or "might makes right." But moral imagination allows for new values to be discovered *in situ,* like an emergent game of catch and the corresponding value of sharing.

are for our eyes only; the very process of contextual philosophical research functions as Deweyan experimental inquiry. Notice that this conclusion, that qualitative work is normative in the philosophical sense, dovetails with the methodological points I made in Chapter 4; Howard Becker, among many others, stresses that any sociological work is normative. By choosing not to take things at face value, refusing to accept the "obvious" state of affairs, the analyst inevitably acquires a political vector. Deweyan pragmatism extends this conception of normative scholarship to explicitly include situated meditations on "the good," as part of any investigation into social worlds. In this way, philosophers of science have a new way to conceive of their work. The normative potential I describe here is not merely theoretical, however. Empirical, qualitative work of this kind is already being done in many neighboring disciplines; there are many tools that can guide our inquiry in a more precise way than "just look at the practice/context." I've collected here a tentative list of new foci that might facilitate experimental inquiry better than Truth or Justice. Most of these tools are borrowed more or less unmodified from science and technology studies. I present them here with special attention paid to the ways in which they are suitable for philosophical attention and adaptation.

Sociotechnical imaginaries

Technoscience would not be possible without its imaginaries, both technoscientific (Marcus, 1995) and sociotechnical (Jasanoff & Kim, 2015). As I illustrated in Chapter 4, successfully deployed imaginaries motivate and sustain collective projects. Most importantly, sociotechnical imaginaries are rich with content that is suitable for philosophical evaluation in a Deweyan mode. Here I will focus on two aspects: desirable futures and boundaries. When a practice is animated by visions of desirable futures (or by feared disasters), we can ask if the future is indeed desirable. What values are represented? For whom is it desirable? In the case of neural engineering, I found myself asking what sort of ideals of the human body and the "body politic" was presupposed by neural prosthetic research. In short, they are ideals of functional legs and an egalitarianism of bodies. Simultaneously, my qualitative research suggests that this ideal is taken to be desirable by neural engineering researchers, but without much input from persons with disability. These features of the imaginary are appropriate objects of philosophical reflection. By making them explicit in new ways and new spaces, the philosopher can start the pragmatic method anew. And given the philosopher's familiarity with a broad (and sometimes competing) range of ethical and political ideals, it should be no problem to identify the sustaining values at work in an imaginary.[7]

The imaginary also encodes the means by which the desirable future is meant to be achieved. This can take the form of institutional or material arrangements, boundaries (asserted or challenged), and attributions of authority or expertise. Recall that members of the Center for Sensorimotor Neural Engineering often emphasized that their project was best achieved by interdisciplinary collaborations, using methods as needed from molecular biology, neuroscience, electrical engineering, with neuroethics present but separate. Non-expert publics, moreover, are not included except as beneficiaries. Simultaneously, ethics is cordoned off as a distinct domain. Philosophers should find a wealth of questions to ask here: how is the production of knowledge appropriate given the desirable future of the imaginary? Is the division of labor indeed as interdisciplinary they say? In more political terms, we can also ask about omissions in the institutional arrangements, like the relative absence of persons with disabilities, or the un-argued for ideal of "egalitarianism of bodies." By starting with imaginaries, rather than just epistemology or ethics, the philosopher can maintain a focus on the practice as it really exists and reveal or re-open unseen possibilities.

Translations

Since the 1980s, Michel Callon and Bruno Latour have advocated for a new symmetry in analysis, bringing objects and materiality into the social relations that define human activity. Latour gives non-human things new standing when he labels them "actants." The resulting framework is a "sociology of associations," a tool with which the analyst traces the paths and interactions of actors without any preconception of what "the social" might be.[8] As a result, the stabilization of matters of fact or phenomena that are characteristic of technoscience can be described in terms of "translation." Callon (1986) famously gives scallops agency in order to illustrate how they were domesticated in St Brieuc Bay. As he tells the story, there are four moments of translation:

> (a) problematisation: the researchers sought to become indispensable to other actors in the drama by defining the nature and the problems of the latter and then suggesting

[7] Jasanoff in (Jasanoff & Kim, 2015) recommends comparison as a reliable way to tease out what is distinct in a given imaginary. Accordingly, philosophers need not rely solely on their intuition to know when they've found an imaginary.
[8] See Latour (2005)

> that these would be resolved if the actors negotiated the 'obligatory passage point' of the researchers' programme of investigation; (b) interessement: a series of processes by which the researchers sought to lock the other actors into the roles that had been proposed for them in that programme; (c) enrolment: a set of strategies in which the researchers sought to define and interrelate the various roles they had allocated to others; (d) mobilisation: a set of methods used by the researchers to ensure that supposed spokesmen for various relevant collectivities were properly able to represent those collectivities.[9]

The four moments give a general account of how the social and natural simultaneously take form, as actors work to gain the compliance of other actors (human or otherwise). And in the spirit of symmetry, Callon treats the interests of human scientists no differently from the the scallops "choice" to anchor on a horsehair substrate rather than nylon.

Successful "translation," as described Callon and Latour, might lead some philosophers to recast the episode in terms of truth or matters of fact. Scallops can be domesticated or they can't, someone might say, and it does not matter how the stakeholders around St. Brieuc Bay came to know this or their motivations to start inquiry. Such a response may be tempting, but it is not a correspondence with reality or the god's-eye view that makes research a success in practice. Technoscientific matters of fact and phenomena are stabilized by their extension beyond the lab and their contextual translation for the host university, the funding agency, top journals, peers, graduate students, and industry partners. My qualitative work in neural engineering revealed some of this complexity, showing how researchers leverage both materials, ideas, and social resources. Philosophers will see truth lurking, of course, in translational processes. But to see truth alone would be to deliberately avoid features of scientific practice and, thus, miss other dimensions of normativity.

The ANT framework is a complementary metric to that of sociotechnical imaginaries, but only if we keep an eye on the values or desirable futures that make translation possible. As Jasanoff stresses in *Dreamscapes of Modernity*, only humans have the capacity to imagine and re-imagine what the world should be like. Accordingly, we must not allow the "flatness" of Latour networks to hide the normativity and power inherent in human action, in translation. It would be too quick, for example, to say that if some technoscientific practice successfully translates than it must have been good technoscience. But there is something noteworthy about successful translation from

[9] Elements of a Sociology of Translation, p1

a pragmatist point of view; a reconfiguration (perhaps many solutions) has been achieved. This achievement, of course, takes the place of many hypothetical solutions or actions that were not pursued. Even after the fact, the philosopher can inquire as to what we left behind on the road not taken.

-Constitutionalisms

Finally, I'd like to mention the lens of constitutionalism, as introduced in Jasanoff (2011). There, Jasanoff encourages us to attend to the ways in which science and technology can shape social orderings, alter our rights and responsibilities, and change what a citizen can reasonably expect from her society. In the other direction, law and legal order can impact knowledge practices and their place in society. As in the previously mentioned tools, the analyst should not presume the primacy of the social or the epistemic, but should be open to a variety of co-productive possibilities. Consider for example the way in which neuroscience shows up in criminal cases, influencing the way we understand and punish the convicted. Denno (2015) for instance shows that the majority of neuroscience in the courtroom is used to mitigate the sentencing of those facing capital punishment, citing a range of diagnoses of brain damage and behavioral disorders. Criminal law, in this way, is shaped both by its own texts and by the book of nature, so to speak. Beyond the realm of criminal justice, we can also find constitutional dynamics in the emergence of new rights holders, new biological citizens united by their genetic or biochemical makeup. Rabeharisoa & Callon (2004) chronicle one such emergence, in which the genetic element of muscular dystrophy becomes part of certain patients' identity. And through the Association of French Muscular Dystrophy Patients, dystrophy researchers find themselves faced with new claimants, patient-activists who assert a right to influence the methods and agenda of biomedical research. These are just a couple of examples of "bioconstitutionalism," but the normative potential of this framework extends broadly to any area of technoscience.

Through lenses of bio-, bit-, or even neuro- constitutionalism, we can inquire what sort of constitutional changes accompany present-day technoscience. In terms of philosophical concerns, there are at least two centers of normative concern identified by Jasanoff. First, there is the issue of care and constitutional protection: "who belongs to communities of moral concern, and who is responsible for taking care of life in those communities?"[10] Unlike the stereotypical property-owning, able-bodied citizen implied in the American Constitution, novel communities and identities created through technoscience demand our legal and moral attention. Second, we can ask also whether the constitution emerging in various applications of technoscience is premature or the product of broad deliberation, as in Jasanoff (2013) "A World of Experts: Science and Global Environmental Constitutionalism." There, she asks whether the global authority of the Intergovernmental Panel on Climate Change, with its narrow definitions of scientific expertise, may be a premature constitutional arrangement. Implicitly technocratic solutions to governance like this one would rarely be subject to broad deliberation were it not for scholars in STS or (eventually) philosophy of science. Along these dimensions, -constitutionalisms articulate a thoroughly normative way to investigate the entanglement of knowledge practices and social order.

Philosophical fault lines in contemporary technoscience

These are a few of the primary theoretical-conceptual lenses from STS that could enable a more symmetrical philosophy of technoscience. They provide the philosopher with objects of study that do not replicate problematic boundary work present in much of philosophy and the natural sciences. They avoid "hypostatization" of truth, justice, or goodness. Doubtless, the selection of tools here is my own, shaped by my interests and fieldwork experiences, and someone else might create a slightly different list. Regardless, the main point is to pick one, and go use it. That's it! My job here is done. "Wait, wait, wait," you are probably thinking. "You said you'd help us understand what good technoscience *is*, and instead you've just given us a method and some unfamiliar theoretical jargon." Fortunately, I can say a bit more about the character of good technoscience in the 21st century. If we abstract away from the STS narratives of Chapter 1 and the qualitative work of Chapter 4, there is a series of questions that emerge. I call these "fault lines in contemporary technoscience," sites where fundamentally philosophical issues are being negotiated whether or not a professional philosopher is around to participate or pay attention. I've found that such fault lines are best revealed empirically, from the bottom-up through sociotechnical imaginaries, ANT analyses, or -constitutionalist lenses, but for my audience of philosophers it is useful to spell out the normative questions more explicitly. As will become evident, the questions will not always be

[10] Jasanoff (2011), p22

easily separable or conceptually independent. Neither are they comprehensive of all the current or future issues associated with technoscience. Nevertheless, these questions can serve to instigate new lines of inquiry within philosophy of science that are better positioned to respond to the challenge of co-production.

Publics and their role in technoscience

What is the public for a given technoscientific practice? Do they have a say in the day-to-day research? Do they have any role in research at all? Should they?

The nominal goal of helping persons with disability and the integration (in some sense) of neuroethics into the CSNE indirectly answered these questions. But there is good reason to think that the answer was not obvious or in any way inevitable; it certainly was not the result of wide deliberation among diverse participants. Some indication comes from scholars who have questioned the role of institutional ethics in protecting technoscientific research from meaningful public interventions. Even the researchers at the CSNE begin to unravel the dominant imaginary, wondering and doubting out loud if neural technology will ever be used by persons with disability (as opposed to the scientific community, able-bodied consumers, or some unknown public). Meanwhile, persons with disabilities rarely play any direct role in neural engineering.

Whether this situation (or any technoscience-public arrangement) is satisfactory is partially an epistemic matter, since the alleged public(s) may or may not be characterized accurately; and even if they do exist as depicted in the imaginary, their voices may not count for anything in the reigning epistemology. Our evaluation of a technoscience-public relation is also a function of our political commitments, since the public(s) may not have the voice they deserve in a democratic society. To fill out my theory of good technoscience in a promissory way: good technoscience will have real publics, not straw publics, and they will be integrated into research in a way that is just and effective.

Novel institutional arrangements, the realization of desirable futures, and responsibility

What are the hopes or fears driving a technoscientific practice? For whom are they motivating? Are they worthy of our efforts?

Imaginaries are particularly useful for opening up these future-oriented aspects of technoscience. My experiences at the CSNE convinced me that the goal of making bodies equal may not be the only (or even a possible) solution to the disabilities associated with spinal cord injury. Comparison with discussions of disability outside of engineering reveals the extreme contingency of the neural imaginary's value-content. However, the dominant narrative driving neural engineering and managing the flow of resources (whether monetary or human) makes resistance difficult for the researchers.

Here is where philosophers of science have a role to play. Even if an imaginary is well-"embedded," re-description by an outsider can create new opportunities to resist the dominant imaginary and its desirable future. The misgivings of neural engineers would have remained an undercurrent were it not for our interviews with them. The point here isn't that neural engineering is "bad" technoscience, but that outside critique can allow for pragmatic reflection and re-alignment on the goals of technoscience. In the ideal case, good technoscience will reflect genuine hopes and fears, ones that are motivating for all relevant[11] stakeholders.

What sort of institutional and organizational arrangements are seen as means to achieving the envisioned desirable future or avoiding the feared outcome? Are they appropriate solutions? What sort of knowledge can we expect from these arrangements?

Some of the normative import of this question can be seen in Galison's treatment of "trading zones," of which he seems to approve implicitly. After all, radar wouldn't have been possible if theoretical physicists and radio engineers hadn't interacted in the right sort of way. At the same time, theoretical physics was in some sense enriched by its encounter with more concrete applications of knowledge. The outcome would have been very different if the scientists had maintained the "pure" character of scientific practice. Of course, Galison primarily justifies the trading zone concept as a way for the historian to understand how actors overcome the problem of incommensurability. Yet, trading zones have an aura of goodness, as productive institutional assemblages. Galison's (2010) work in the volume *Trading Zones and Interactional Expertise*, along with Michael Gorman and others, makes it explicit that the trading zone is a practically useful tool, which we can apply to institutional design in the present day. My point here is that the network of technoscience, both "internal" and "external" matters and can be the object of explicit evaluative reflection. We should critically assess how knowledge and things, products of technoscience, travel

[11] See the question of publics

beyond labs and academic departments.

The contingent and overlapping social worlds that connect technoscience to society shape how its products are "translated". Just think: wouldn't neural engineering create very different knowledge and technologies if it were as tightly connected to disability activist groups as it is to the medical devices industry? Climate science, too, says different (more or less alarming) things about humanity's future depending on whether it is conducted in a university department or in a consensus-oriented, interdisciplinary panel. The specific socio-political procedure by which technoscientists "raise the world,"[12] affects what we can know and what we can do. And while these arrangements are often not the plan of any one person or collective, philosophers of science and social epistemologists can make them objects of study, questioning their effectiveness and fairness. To return to the ideal, good technoscience will utilize the institutional structures, expertise, and methods that are fitting for the desirable futures that sustain it and for the publics it serves.

How is authority and accountability distributed within the network of practice? Does it allow for fair attributions of responsibility or does it create a culture of "organized irresponsibility?"[13]

The sociologist Ulrich Beck fears that modern-day society is characterized by "organized irresponsibility," an unfortunate result of the distributed character of large, complex institutions like technoscience and international trade.[14] What one person or collective could we possibly blame for the existence of nuclear waste, climate change, or the Internet? A smaller scale version of this worry appeared in my discussions with neural engineers; each researcher described themselves as part of a large network that relieved them of most control (and for some, responsibility) over their technological contributions to society. Beck says that since we can't get attributions of responsibility to stick to individuals in technoscience, we blame politicians instead.[15]

Given that publics do exist and the technoscience is not conducted in a galaxy far, far away from humanity, the social organization ("real" or in imaginaries) has profound effects on our ability to demand accountability from technoscientific actors. If there is no one who has the task of, say, reaching out to the publics of neural engineering, then we will have no one to reasonably blame for their absence in the CSNE. Along similar lines, particular sociotechnical imaginaries can

[12] to steal Latour's phrase
[13] a la Ulrich Beck's "Risk Society"
[14] See Beck (1998).
[15] Beck (1998), p14

rule out the possibility of technoscientific accountability, including Kuhn's internalist description of scientific communities and my interviewee's ambiguation of neural engineering's "end user." Philosophers of science should ask how social structure and notions of responsibility are tied in technoscience. In this spirit, one last ideal: good technoscience is organized to enhance rather than undermine public accountability of the practice.

Good technoscience: for philosophers and everyone else

Questions do not make for a very concise account, especially when they come in a long series. But it is possible to work backwards from the fault lines I identify and to reconstruct some ideals. If we allow ourselves to hypostatize just for a moment, my tentative account of good technoscience looks something like the following: Good technoscience *i*) is oriented towards real publics and integrates their voices in a just way, *ii*) reflects genuine hopes or fears that are motivating for relevant stakeholders, *iii*) leverages the right institutions, experts, and methods to effectively achieve desirable futures (or avoid disasters), *iv*) and is organized to enhance rather than undermine public accountability of the practice. A pragmatist account of good technoscience, as this list shows, cannot provide much in the way of useful abstract prescriptions. Dewey would probably laugh at my naive attempt. The definition of "good" technoscience can only be made (co-produced) in the moment, with whomever and whatever is present. For this reason, the questions listed above are a better place to start. I hope, nevertheless, that this sketch hints at the normative concerns that will be raised in qualitatively-grounded, contextual and co-productionist philosophy of science. Philosophers' paranoia about losing their normative edge is just not applicable here.

Perhaps to the dismay of some readers, my account of good technoscience provides very little in terms of universalizable principles or ideals. Yet, what it lacks in content it makes up for in spelling out a particular method and mode of being for philosophers of (techno)science. It entails that we inhabit new spaces and engage in a form of critique that resembles the constructivist work being done elsewhere in the STS sphere. Since science and technology provides tacit answers to moral, political, and epistemic questions, there are always issues of philosophical consequence waiting to be addressed. The key is to keep this fact in mind—STS tools can help—and not lose oneself in the archives, be mesmerized by formula and lab equipment, or collapse into a causal-explanatory way of thinking. The goal of co-productionist critique is never to unpack things for unpacking's sake, to deconstruct just to deconstruct. Without a sensitivity to the normative questions that interest us even as analysts, constructivist scholarship really will "run out of steam," as exhibited by Latour's (2004) moment of self-doubt. By revealing the tacit normative content embedded in technoscience, we can reflect more directly on whether they reflect our values as a society or those of intended beneficiaries. This insight brings some clarity to the questions that ended Chapter 2; Longino's suggestion that she somehow represents a community beyond herself could be read as this type of critique. Philosophy of science could thus be oriented towards fulfilling Jasanoff's imperative for STS critique, always asking "Why is it so, and must it be?"[16]

Though philosophers lose their authority to stipulate the conditions of well-ordered technoscience, my account of good technoscience opens up a range of new engagements to improve science and engineering. Philosophers can join a community of like-minded critical scholars of science and technology, working together on shared questions of normative import. By reconnecting philosophy of science to STS in a robust way, we bring together the two paths that parted ways after Kuhn's *Structure*. At the same time, my account pushes us to rethink how we intervene on technoscience itself. Embedded social scientists, anthropologists, and ethicists are obligated to resist the increasingly common sentiment that they are somehow experts in matters of "responsible innovation" and research ethics; they have to undo the professionalization that has already begun in the case of bioethics, instigated in part by the actions of funding agencies. Deliberating on the nature of good technoscience, as a co-productive process, is as open to mail couriers and lab technicians as it to professional academics. The next and final chapter will explore this humbling thought.

References

Alexander, T. M. (1993). John dewey and the moral imagination: Beyond Putnam and Rorty toward a postmodern ethics. *Transactions of the Charles S. Peirce Society*, *29*(3), 369–400.

Beck, U. (1998). The politics of risk society. In J. Franklin (Ed.), *The Politics of Risk Society*. Cambridge: Polity Press.

Callon, M. (1986). Some elements of a sociology of translation. In J. Law (Ed.), *Power, Action, and Belief* (pp. 1–40). London: Routledge.

Charlton, J. I. (1998). *Nothing about Us Without Us: Disability Oppression and Empowerment*. Berkeley: University of California Press.

[16]in "The Floating Ampersand: STS Past and STS to Come", forthcoming

Denno, D. W. (2015). The myth of the double-edged sword: An empirical study of neuroscience evidence in criminal cases. *BCL Rev*, *56*(2), 493–551.

Dewey, J. (1922). Valuation and experimental knowledge. *The Philosophical Review*, *31*(4), 325–351.

——— (1938). *Logic: The Theory Of Inquiry*. New York: Henry Holt and Co.

Galison, P. (2010). Trading with the enemy. In M. E. Gorman (Ed.), *Trading zones and Interactional expertise: Creating New Kinds of Collaboration*. Cambridge MA: MIT Press.

Jasanoff, S. (Ed.) (2011). *Reframing Rights: Bioconstitutionalism in the Genetic Age*. Cambridge MA: MIT Press.

Jasanoff, S. (2013). A world of experts: Science and global environmental constitutionalism. *BC Envtl Aff L Rev*, *40*(2), 439–452.

——— (forthcoming). The floating ampersand: STS past and STS to come. *Engaging Science, Technology, and Society*.

Jasanoff, S. & Kim, S.-H. (2015). *Dreamscapes of Modernity: Sociotechnical Imaginaries and the Fabrication of Power*. Chicago: University of Chicago Press.

Latour, B. (2004). Why has critique run out of steam? from matters of fact to matters of concern. *Critical Inquiry*, *30*(2), 225–248.

——— (2005). *Reassembling the Social: an Introduction to Actor-Network Theory*. Oxford: Oxford University Press.

Marcus, G. E. (1995). *Technoscientific Imaginaries: Conversations, Profiles, and Memoirs*, volume 2. Chicago: University of Chicago Press.

Rabeharisoa, V. & Callon, M. (2004). Patients and scientists in french muscular dystrophy research. In *States of Knowledge. The Co-Production of Science and Social Order* (pp. 142–160). London: Routledge.

6. Afterword - Interventions by an Embedded Co-Productionist

"Why?" questions, destroyers of worlds

At this very moment, an astrophysicist somewhere is sitting down at the computer to draft a conference presentation. She thinks she has found a habitable exoplanet, far beyond the limits of our solar system and invisible to the naked eye. The strange thing about this type of astronomy is that no one ever really sees the exoplanet. We have to infer the planet's existence and its properties from the way it dims the star that it orbits. Visually, the effect is extremely underwhelming. Yet, based on some mathematical modeling and some pre-determined statistical threshold, we can say that planet Earth has a twin somewhere out there. The details of the modeling process are not so important here, but I have always been fascinated by the way that another Earth, perhaps filled with life not unlike our own, can pop in and out of existence, depending on whether the astronomer is held to a 3-sigma standard of significance or a 5-sigma standard. Strangely enough, this sort of statistical cosmogony actually happens when competing astrophysics laboratories challenge one another's findings. Stranger still, this scenario was the starting point for my dissertation project.

I suffered an unhealthy obsession with exoplanets during Professor Alison Wylie's graduate seminar (Science and Values) as we were discussing the recent spate of "embedded" social science and "adjacent" ethics foisted upon scientists and engineers. These interventions are typically accompanied by justifications in terms of "responsible innovation" and ethical science. Sociologists or philosophers can help ensure that the right values are represented in technoscience, or so the argument goes. In the United States, this approach is best exemplified by Erik Fisher and Dave Guston's STIR (Socio-Technical Integration Research) Project. Their idea was to use funding from the National Science Foundation to place "embedded humanists"—these particular humanists happened to be graduate students—in ten labs around the world. Once embedded, the humanist's job was to ask Socratic "why?" questions at "decision points" in the research. One can only imagine the gadfly-like effect of such behavior.

Surprisingly, Fisher and Guston come to describe their approach not as inhibitory but as productive. Fisher first promoted this idea in a (2010) correspondence piece in *Nature*, where he claims that "midstream" reflection brings neglected socio-ethical values back into the research context. But interestingly, he also pitches the idea that the STIR protocol is useful *even by the internal standards of scientists and engineers*. There, he hints provocatively: "reflections on responsible innovation generated novel ideas for antenna structures and nanoparticle synthesis for researchers at ASU's Center for Single Molecule Biophysics." A year later, one of the STIR participants explicitly acknowledges both the socio-ethical and the epistemic side of reflection: "Second-order reflective learning involves reflection 'on' the research system, including the value-based socio-ethical premises that drive research, the methodological norms of the research culture, and the epistemological and ontological assumptions upon which science is founded.[1]" What were the results of such self-questioning? Schuurbiers reports that his repeated prodding of researchers led to discussions of deeply philosophical topics, including reductionism and the underdetermination of theory by data. But Schuurbiers says no more about how these discussions affected the more epistemic parts of practice, if at all.

This overall theme of scientist-approved ethics engagement is repeated in the three yearly NSF reports compiled by Fisher and Guston. In year one, they highlight STIR activities as creating socially responsive scientists via the cultivation of reflective habits:

> Typically, the "considerations" (values, goals, concerns and other decision criteria) that practitioners articulate while describing their everyday research decisions increase in number and broaden in quality as a result of the routine collaborative exercises. Values mentioned can expand to include ethical, environmental, public and other societal dimensions of the research in question.[2]

The authors claim this expansion led to real changes in both material practices and institutional norms. Fisher and Guston resubmitted much of this content for the second yearly report, but with a few notable additions. First, they assert that even scientists value reflective learning. They describe how, in one lab, scientists decorated their wall with a diagram of positivism and constructivism, a way to remind themselves of different "ways of thinking."[3] In another case, the embedded humanist was praised for pointing out implicit contradictions in the lab's research vision. But despite these positive instances, Fisher and Guston issue a qualified claim: "Changes

[1]Schuurbiers (2011)

[2]Fisher, E. and Guston, D. (2010) STIR: socio-technical integration research (NSF#: 0849101), annual report for 04/2009-03/2010. Retrieved from: http://cns.asu.edu/resource/13420

[3]Fisher, E. and Guston, D. (2011) Annual report: February 25, 2011. See p4. Retrieved from: http://cns.asu.edu/resource/13421

in practice, when they do occur, are generally seen as productive by the laboratory practitioners." They mention only one case where a researcher called the STIR methods a "waste of time."

Less than a year later in 2012, Fisher and his colleagues begin to use stronger language to assert the compatibility of STIR methods with scientists' and engineers' internal goals. In the year-three report to the NSF, they summarize the results of the study: "This study provides preliminary evidence that such activities enable laboratory work to become more sensitive to its potential societal implications, without compromising laboratory research, education or strategic goals."[4] And by 2013, the language is stronger than ever. In an edited volume on responsible innovation, Fisher and Arie Rip (2013) write that STIR "has been shown to broaden and enhance R&D decision-making processes." And this enhancement, allegedly, is one that reunites the value sets of science and society.

Returning to the issue of twin Earths, reading about the STIR project made me question the idea of scientist-approved ethics engagement. I found it slightly misleading for the authors to imply that one can tack on socio-ethical values to a practice and preserve, untouched, some core notion of scientific or epistemic success. Taking Fisher and Guston at their word that values were indeed changing in a substantive way, I suspected that there was something more transformational at work, and exoplanet research came to mind. An embedded humanist in an astrophysics lab, asking the right "why?" questions, might have the power to bring worlds into existence and strengthen visions of a galactic humanity or, at the other extreme, unceremoniously squash those hopes as a gambler's fantasy. It's the difference between tentatively reporting a probability of 95.45% and literally painting a picture of a vast alien sea and green continents, which NASA-employed artists are keen to do. It was this silly thought experiment that made me wonder what we academics are doing when we intervene on science and engineering. Is it an epistemological exercise or is it to "make science ethical" or both? If the latter, with what authority or expertise do embedded humanists remake technoscience?

These questions gained a personal significance when I became a neuroethics fellow at the Center for Sensorimotor Neural Engineering (CSNE), tasked to engage with neural engineers about the ethics of their work. What was once a outlandish (extraterrestrial, even) thought experiment became a central concern of my day-to-day life. Suddenly, I was faced with people—some were

[4] Fisher, E. and Guston, D. (2012) Annual report: January 1, 2012. See p2. Retrieved from http://www.cspo.org/library/reports

senior researchers—who labeled me as "ethicist" and would sincerely ask me for advice. Others would give me responsibilities when we talked about the difficulties of making neural engineering ethical: "that's why we have you," a friendly sentiment which nevertheless made me nervous. Anxiety aside, I still managed to form my own opinions about what is right and wrong with brain-computer interfaces. Thinking back to my coursework in disability studies, I was concerned about the CSNE's persistent need to "fix" people's bodies, even though persons with the targeted disabilities are few and far between in neural engineering contexts. Like the embedded humanist in an astrophysics lab, I felt that I had the power to critique the underlying vision of the research. Unfortunately, I lacked any vocabulary or theory to understand what I was doing. Most of my training in philosophy of science and was completely ineffective in this respect, since it presumed that ethics and social order were things to be evaluated independently of scientific practice. Yet, here I was trying to figure out why neuroscientists and engineers cared so much about connecting brains to the outside world in a mess of cables and flesh.

Interventions for ethics! Or for knowledge?

Without a co-productionist perspective, the idea of the "embedded humanist" can be justified in a few different ways. Sometimes, the stated motivation is purely ethical, underplaying the epistemic. My own "Neuroethics Thrust" at the CSNE is one such effort; we are called "neuroethicists" and focus on topics like identity and moral responsibility.[5] Mildred Cho and her team at Stanford present another, the "Stanford Benchside Model: Research Ethics Consultation." Cho et al. (2008) adapting a theme from medical ethics, treat society as a sort of patient in need of care. They assume "the overall goal of maximizing the benefits and minimizing potential harms of research to society." The authors elaborate this goal a bit by including various stakeholders as part of society, including the researcher, research subjects, as well as the institution and the general public. But notably absent is any goal related to the epistemic content of the research itself. Even though they say that their activities are analogous to biostatistics consultations, they never make the connection that the ethical and epistemic could be co-produced. At the other extreme, Helen Longino and other value-sensitive epistemologists assert that epistemic acceptability requires the right mix of criticism and diversity in the research community. So an embedded humanist

[5]This focus is clear in our article: Klein et al. (2015)

might improve knowledge practices, but the ethico-political effects are consequently subsumed as a variable in the knowledge equation.

The idiom of co-production suggests that these one-sided justifications via ethics or epistemology are misleading. Of course, some would-be "embedded humanists" have noticed this. Philosopher Inmaculada de Melo-Martin asserts that the epistemic and the non-epistemic are intertwined in scientific contexts. In her (2009) article "Creating Reflective Spaces", she justifies her assertion by citing several case studies which illustrate the interplay between epistemic and ethical and social values in biomedical contexts. Worrall (2008), for example, focuses on the primary ethical concerns present in randomized clinical trials: informed consent, the timing of trial completion, and the ethical status of the trial itself. He argues that these three issues cannot be addressed without first understanding the epistemological aspects of the study at hand. The ethical justification for a randomized trial, for instance, requires a genuine lack of certainty about which treatment is most effective. But referring to a lack of certainty, he argues, implies a whole set of evidential and doxastic definitions that require an explicit formulation. De Melo-Martin uses cases like this one to emphasize the underexamined connections between the epistemic and the ethical. Science and medicine would benefit, she suggests, from a concerted effort to unearth and understand these connections, even "upstream" in the research pathway, in labs and clinics, well before the typical stages of ethical analysis and government regulation.

And whereas Longino does not specify who would best fulfill the critical role adjacent to science, de Melo-Martin argues that philosophers are particularly well-suited for the task. She argues for this in two ways. First, she lists traditional virtues that philosophers are supposed to have: the ability to bring a normative perspective, the ability to pick out fallacies, as well as extensive training in exorcising unacknowledged background beliefs and hidden assumptions. These characteristics, she suggests, enable philosophers to "alert scientists to ethical and epistemological elements in their work, their presumed assumptions, and the values underlying their investigations." Second, de Melo-Martin cites several instances where philosophers have provided both ethical *and* epistemological "examinations" of biomedical science, in addition to revealing science as a situated social practice (including one study of her own).[6] These examples form the premises of a loose sort of induction; de Melo-Martin hints that philosophical interventions "might well lead to better

6

science."[7] Thinking back to the STIR project reports, I daresay Fisher and his colleagues would concur. In both cases, the reader is supposed to believe that science will be improved rather than inhibited, as ethical and epistemic assumptions are unearthed.

"Good technoscience" in the midstream

De Melo Martin and Fisher and Guston all acknowledge that intervening on technoscience at the "midstream" has epistemic and ethico-political components, and they are advocating for a particular arrangement of knowledge production that just happens to require their presence (or the presence of people like them) in labs all over the world. Their standards of "better" science cannot be separated from the desirable future that they envision (if only implicitly). I suggest we see them as strategic co-productionists, or perhaps Skinner box co-productionists, since they may not think of themselves precisely in that way. Accordingly, their work is just as amenable to analysis by sociotechnical imaginaries or -constitutionalisms as the neural engineering in Chapter 4. We can ask why particular interventions on technoscience are framed the way they are. What is the desirable future and how is it to be achieved? What new rights, responsibilities, and identities emerge when national funding agencies pay to place philosophers and other humanists in labs? Answering these questions requires that we not limit ourselves to studying the strategic justifications of would-be interveners, as when Fisher and Guston are justifying their work to the National Science Foundation. As I did with Longino's epistemology, we must read between the lines of official texts and op-ed pieces and make explicit the full vision for well-ordered technoscience. Ethnography and participant observation is probably necessary too, watching the watchers, so to speak.

This realization pushes me to reflect more aggressively on my work as an Neuroethics fellow and as a co-productionist philosopher of science. Each role tacitly answers the question of how technoscience should be ordered in present day America. As neuroethics fellow, it is obvious to me now that I have been enrolled by an imaginary that is not of my invention. The idea that a graduate student and a few of his colleague friends can be entrusted to make neural engineering ethical draws on and perpetuates the imaginary of elite bioethics committees, who stand in for the public. Jasanoff (2011) describes this exact dynamic as a "constitutional moment in govern-

[7] "Creating reflective spaces", see p43

ing science and technology," in which democracy itself is being defined:

> One sees here little of the unruliness and raw contestation of democratic debate. Rather, ethics committees engage in a polite process of opinion formation, oriented toward elite consensus-building, in which the values and sensibilities of a very few, highly educated, articulate individuals stand in for the untrained, and allegedly uninformed, preferences of the multitude. A subtle transformation occurs in the dynamics of deliberation when research is subjected to ethical reflection rather than political debate. Ethical expertise imperceptibly gets subordinated to technical expertise, so that those with privileged understanding of the scientific subject matter under discussion come to be seen, and also to see themselves, as the best representatives of the public's moral commitments.

Whether or not this is a desirable consequence, Jasanoff's emphasis on democracy shows that the stakes are higher than de Melo Martin or Fisher and Guston suggest. What's missing in justifications of midstream engagement is a more direct consideration of how "reflective spaces" and "responsible innovation" open technoscience to some, but close it down for others. Consider the fact that I've repeatedly typed the slogan "nothing about us, without us," but I have never knowingly collaborated with someone who identifies as a person with a disability; it is my own able body that attends lab meetings and types up neuroethics articles.

If we agree with Castoriadis' suggestion that there is not such thing as political expertise, only political judgment, it seems likely that the midstream engagement literature has taken a wrong turn. The question of engagement may not amenable to debate in esoteric academic journals or at closed door workshops. And neither should we take embedded neuroethics at face value. Fortunately, I am free to imagine broader and better forms of engagement. In this very dissertation, Chapter 5 could be read as one such proposal. My pragmatist framework allows me to think of my presence in technoscientific spaces as a form of experimental philosophy. In this role, I do not pretend that I am an adequate representative of the public or of persons with disabilities. Though I may have them in mind in some of my engagements, they are always filtered through my own sensitivities and blindspots. Unlike a democratic representative, I can never be recalled by my publics if they are unhappy. Embedded co-productionist philosophy is not representation. It is participation. Although I may have some of the philosophical skills and dispositions that are highlighted by de Melo Martin, I see myself primarily as one more person at the table, co-producing value and knowledge in technoscientific spaces. Skill, to my mind, should not trump equality.

My scholarly self-understanding will strike some readers as egocentric, as a reduction of "midstream" engagement to personal interest. And that would make for a quick end to this topic. Nothing redirects philosophical attention quicker than an admission of vulgar selfishness. However, it would be dishonest to say there is no desirable future or vision of an orderly society underlying the idea of constructivist critique. In my own distillation of STS critique and pragmatist philosophy, there is a persistent optimism towards the action of human imagination and reason. Co-productionist scholarship demands forms of life in which we continually ask, "What world do we want to live in?" never satisfied with yesterday's normative conclusions. It projects a future in which we never sleepwalk through technoscientific change or allow the continual re-definition of democracy to happen without us. The foundational premise is that the world is better when we exercise our uniquely human capacities and engage whole-heartedly with normative questions. In this picture of society, co-productionist philosophers of science are neither experts nor rulers. Their scholarly critique is nothing more than a citizen's duty to imagine and re-imagine how things should be.

References

Cho, M. K., Tobin, S. L., Greely, H. T., McCormick, J., Boyce, A., and Magnus, D. (2008). Strangers at the benchside: research ethics consultation. *American Journal of Bioethics*, 8(3):4–13.

de Melo-Martín, I. (2009). Creating reflective spaces: interactions between philosophers and biomedical scientists. *Perspectives in biology and medicine*, 52(1):39–47.

de Melo-Martín, I. and Intemann, K. K. (2007). Can ethical reasoning contribute to better epidemiology? A case study in research on racial health disparities. *European Journal of Epidemiology*, 22(4):215–221.

Fisher, E., Biggs, S., Lindsay, S., and Zhao, J. (2010). Research thrives on integration of natural and social sciences. *Nature*, 463(7284):1018–1018.

Fisher, E. and Rip, A. (2013). Responsible innovation: multi-level dynamics and soft intervention practices. In Owen, R., Bessant, J., and Heintz, M., editors, *Responsible Innovation: Managing the Responsible Emergence of Science and Innovation in Society*, pages 165–183. Wiley, Hoboken.

Jasanoff, S. (2011). Constitutional moments in governing science and technology. *Science and Engineering Ethics*, 17(4):621–638.

Klein, E., Brown, T., Sample, M., Truitt, A. R., and Goering, S. (2015). Engineering the brain: Ethical issues and the introduction of neural devices. *Hastings Center Report*, 45(6):26–35.

Schuurbiers, D. (2011). What happens in the lab: Applying midstream modulation to enhance critical reflection in the laboratory. *Science and Engineering Ethics*, 17:769–788.

Worrall, J. (2008). Evidence and ethics in medicine. *Perspectives in Biology and Medicine*, 51(3):418–431.