



Linking Mind to Molecular Pathways: The Role of Experiment Tools

John Bickle^{1,2}

Received: 20 May 2019 / Accepted: 1 June 2019 / Published online: 8 June 2019
© Springer Nature B.V. 2019

Abstract

Neurobiologists talk of linking mind to molecular dynamics in and between neurons. Such talk is dismissed by cognitive scientists, including many cognitive neuroscientists, due to the number of “levels” that separate behaviors from these molecular events. In this paper I explain what neurobiologists mean by such claims by describing the kinds of experiment tools that have forged these linkages, directly on lab benches. I here focus on one of these tools, gene targeting techniques, brought into behavioral neuroscience from developmental biology more than a quarter-century ago. Discussion of this tool does more than illuminate these claims by neurobiologists, however. An account of its development shows the doubly dependent role that theory plays in neurobiology. Our best current theories about “how the brain works” depend entirely on the experiment tools neuroscientists have available. And these tools get developed via the solution of engineering problems, not the application of theory. Theory is thus of tertiary importance in neuroscience, not of the primary importance that many cognitive scientists assume it to occupy.

Keywords Mind-to-molecular-dynamics linkages · Theory-centrism in neuroscience · Experiment tools · Gene targeting techniques

1 Linking Mind to Molecular Dynamics

In the introductory chapter to the Fourth Edition (since superseded by a Fifth Edition) of their monumental *Principles of Neural Science*, neurobiologists Eric Kandel, James Schwartz, and Thomas Jessell offer a remarkable assessment of what neural science provides:

✉ John Bickle
jbickle@philrel.msstate.edu

¹ Departments of Philosophy and Psychology, Mississippi State University, Starkville, MS, USA

² Department of Neurobiology and Anatomical Sciences, University of Mississippi Medical Center, Jackson, MS, USA

This book ... describes how neural science is attempting to link molecules to mind—how proteins responsible for the activities of individual nerve cells are related to the complexities of neural processes. Today it is possible to link the molecular dynamics of individual nerve cells to representations of perceptual and motor acts in the brain and to relate these internal mechanisms to observable behavior. (2001, 3–4).

Their claim is doubly remarkable when one realizes that ‘today’ refers to 2 decades ago.

Cognitive scientists, including many cognitive neuroscientists, will dismiss this claim as so much hubris. So many levels of so many kinds—theories, phenomena, targets of distinct sciences, mechanisms—lie between the behaviors indicative of higher cognitive functions and “the molecular dynamics of individual nerve cells”! How do neurobiologists purport to have established such “linkages”—and for 20 years, no less?! The answer I am going to suggest on neurobiologists’ behalf is: through the use of novel experiment tools, mostly developed over the past 3 decades, which permit new experimental designs. In short, these tools permit the direct testing of such mind-to-molecular-dynamics linkages “in a single bound,” right on the laboratory bench.

In a short paper aimed at such a huge point, one must choose arguments carefully. I am going to focus on one experiment tool that was especially prominent in the work Kandel et al. allude to: the use of gene targeting techniques in behaving mammal model organisms. However, the story of this one tool does more than show how neurobiologists understand these mind-to-molecules linkages. Attention to how this tool developed, and became popular in neuroscience practice, also points to a much more limited role for theory than many envision. Rather than being the crux point on which everything else depends, and thus whose perceived lack generates a desperate desideratum for current efforts, theory turns out to be doubly dependent, and hence of tertiary, not primary, importance. Our best confirmed theory is totally dependent on what our experiment tools allow us to manipulate. And those tools developed by way of solving engineering problems, not by applying theory.

The background philosophy of science at work here comes from Hacking’s (1983) (self-described) “little” book, especially the chapters in Part II where he sought to spur a “back to Bacon” movement. Philosophers of neuroscience have focused lately on experimentation (Bickle 2003; Craver 2007; Sullivan 2009, 2010; Silva et al. 2014; Bickle and Kostko 2018). This includes the role, use, and development of experiment tools (Bickle 2016, 2018; Robins 2016, 2018; Sullivan 2018). However, a point Hacking made about microscopes has not been appreciated about neurobiology’s tools: “Theory has only a modest amount to do with building these ingenious devices. The theory involved is mostly of the sort we learn in Physics I at college. *It is engineering that counts*” (1983, 199; my emphasis). Substitute ‘undergraduate molecular biology’ for ‘Physics I’ and Hacking’s quote carries over directly to the tools that were used to discover Kandel et al.’s mind-to-molecular-dynamics linkages. So in addition to showing cognitive (neuro-)scientists and philosophers what neurobiologists mean when they speak of such linkages, I also hope to disabuse them of their theory-centrism.

2 Theory-Centrism in and About Neuroscience

I start with theory-centrism, encapsulated in the appeal for “more theory” in neuroscience. This appeal is not new. In 1979, while he was still a self-admitted “novice” in neuroscience, Francis Crick wrote an invited commentary on a series of publications in *Scientific American* presenting then-state-of-the-art neuroscience, written by the field’s leaders. He remarked that these articles “gave a good general idea of the progress that has been made” toward understanding the sensing, cognizing, emoting, and action-guiding brain. Still,

what is conspicuously lacking is a good framework of ideas within which to interpret all these different approaches. Biochemistry and genetics were in such a state until the revolution in molecular biology. It is not that neurobiologists do not have some general concept of what is going on. The trouble is that the concept is not precisely formulated ...How then should a *general theory of the brain* be constructed? (1979, 133; my emphasis)

As Crick’s involvement in professional neuroscience deepened, he maintained this early judgement, that the field both lacks and desperately needs general theory. In one of his last scientific publications, co-authored with Christof Koch, the pair offer “a framework ... for explaining the neural correlates of consciousness in terms of competing cellular assemblies” (2003, 119). They are careful to point out that their ‘framework’ is not yet a theory: “a framework is not a detailed hypothesis or set of hypotheses; rather, it is a suggested point of view for an attack on a scientific problem” (2003, 119). A ‘framework’s’ epistemological status and offerings fall short of what a good theory provides: “A good framework is one that sounds reasonably plausible relative to available scientific data and that turns out to be largely correct. It is unlikely to be correct in all the details. A framework often contains unstated (and often unrecognized) assumptions, but this is unavoidable” (2003, 119). So after nearly a quarter-century of working in neuroscience, one of the foremost scientist of the twentieth century went to his grave insisting that, for visual consciousness at least, neuroscience was still in the prolegomenon stage with regard to theory, and the science was worse off for being so.

Crick’s lament was not his alone. Neurophilosopher Patricia Churchland, in her book initiating the field, likewise insisted that, at least concerning how ensembles of neurons work to generate complex behaviors and cognition, “there is no widely accepted theoretical framework, nor even a well-defined conception of what a theory to explain such things as sensorimotor control or perception or memory should look like” (1986, 403). And while she admitted “some sympathy” for those neuroscientists who judge that “theorizing about brain function is ... slightly disreputable and anyhow a waste of time,” nevertheless she offers a number of reasons in support of “the value of theory” (1986, 403–407). Chapter 10, nearly 80 of the book’s 482 pages, presents three potential “theories of brain function,” each one “illustrat[ing] some important aspect of the problem of theory in neuroscience” (1986, 411). This same attitude carried over a few years

later, into her computational neuroscience primer co-authored with prominent neuroscientist Terence Sejnowski: ““Data rich but theory poor” is a description frequently applied to neuroscience,” and “in one obvious respect, this remains true, inasmuch as we do not yet know how to explain how brains see, learn and take action” (1992, 16). The hope they express throughout the book is that then-newly-emerging neurocomputational models will provide our best strategy for generating missing but badly needed brain theory.

Fast forward into the twenty-first century, and this beat goes on. A decade ago philosophers of neuroscience Ian Gold and Adina Roskies sounded the familiar lament: “Neuroscience ... has very few broad theories” (2008, 351). Invoking the term used by Crick and Koch, they insist that “the field is governed by a few frameworks—a crude physicalism and perhaps computationalism—but these serve as fundamental or grounding assumptions rather than theories” (2008, 351). Frameworks differ from theories by lacking the latter’s desired epistemic features: “they don’t provide neuroscientists with predictive powers in the way that physical theories do” (2008, 351). And this lack distinguishes neuroscience from sciences that benefit from having developed, overarching empirically confirmed theories about their target phenomena: “Given this lack of theoretical richness, and the rather local character of the theories that do exist, neuroscience looks quite different from both physics and evolutionary biology” (2008, 351). As recently as a few years ago, Churchland and Sejnowski still lament: “neuroscience is theory-poor” (2016, 667).

One motive behind my recent focus on tool development experiments in neurobiology has been to challenge this rampant theory-centrism. Neuroscience, I claim (Bickle 2015, 2016, 2018) has lots of good, well-confirmed theory; the suggestion that the brain is a scientific mystery is simply false. But the good, well-confirmed theory neuroscience has—the mechanisms of neuronal conductance and transmission at chemical synapses, receptor and ion channel function, mechanisms of synaptic plasticity, including its molecular-genetic and epigenetic mechanisms, details of anatomical circuitries linking neurons to other neurons, and ultimately to sensory receptors and muscle tissue—all resulted directly from the development and ingenious uses of experiment tools. So the idea of fruitfully theorizing about “how the brain works” independently of developing and using new experiment tools is totally alien to neuroscience’s most successful practices. My recent focus on experiment tools is consonant, therefore, with the more widespread science-in-practice approach that has gained traction in recent philosophy of science (<http://www.philosophy-science-practice.org/>).

This narrative begins in Bickle (2015), which describes new designs for intervention experiments made possible by new research tools that had entered neuroscience in the 1980s and early 1990s. These tools included stimulating microelectrodes with tip dimensions small enough to be inserted into cortical microcolumns of similarly tuned neurons in non-human primate brains; receptor- and even protein subunit-specific pharmacological agonists and antagonists; and gene targeting techniques. No longer were neurobiologists limited to measuring individual neuron responses to complex stimuli or during complex behaviors. Now they could intervene, with increasing precision, into cellular or molecular processes correlated with specific cognitive functions, in behaving animal models, and test these correlations as

causal-mechanistic explanations, directly on the lab bench. These experiment tools and designs thus gave neuroscientists means to answer a criticism raised by Marr (1982, chapter 1): that neuroscience could not *explain* cognition, but was instead a purely descriptive enterprise, and so had to be supplanted, on Marr's view with functional and algorithmic explanations provided by cognitive scientists. By the early 1990s, data gathered in intervention experiments in behaving animal models countered Marr's criticism. Neurobiological explanations of cognitive functions were now genuinely explanatory—causal-mechanistically so. And although this point has been forgotten by many, Marr's entire three-level approach to cognitive science was built upon his criticism that neuroscience was descriptive, not explanatory.

In Bickle (2016) I turn attention more explicitly to how these new experiment tools were developed. I investigate two influential tools. The first was gene targeting techniques. The second was optogenetics and chemogenetics, by which expressible genes for entire receptors sensitive to light stimuli or nonbiological pharmacological agents are inserted into the DNA of specific neurons. When expressed, these receptors offer experimenters unprecedented control over the activation or silencing of these specific neurons in behaving animals.

My (2016) metascientific analysis of the development of each of these tools traces them back to their *motivating problems*; and to their *initial* and *second-phase hook experiments* by which they first captured the attention of specialists, then of scientists more generally and sometimes even of the general public. My analysis culminates in a very different account of what drives scientific revolutions in neurobiology than Kuhn's (1962) famous "paradigm replacement" account. I find that the development of specific new experiment tools is the principal driver behind revolutions in the field, at least of ones recognized as such by practicing neurobiologists (which are not always the ones recognized by historians, philosophers, and sociologists of science!). The development of new experiment tools is only one small component of a Kuhnian "paradigm." And the actual development of these new tools likewise typically has nothing to do with Kuhnian anomalies, crisis science, or the development of alternate paradigms, the principal drivers of revolutions according to Kuhn.

I challenge theory-centrism even more explicitly in Bickle (2018). Working specifically with the development of optogenetics. I there point out how features of their development illuminate all the premises of Ian Hacking's (1983) famous "microscope" argument for the relative independence of "the life of experiment" from theory. Hacking stressed features such as laboratory tinkering and "fooling around"; mistaken initial theory-based judgments about the "impossibility" of developments like the electron microscope; the experiment-in-practice origins of some key features of microscopes; the pedestrian level of theory involved in microscope function; and the confidence scientists retained about the veracity of light-microscopic images despite vast changes that occurred in the background theoretical understanding of the physical phenomenon producing them (from an absorption to a diffraction phenomenon). Every one of these "life of experiment" features has a direct counterpart in the more recent development of optogenetics. To the challenge that "a lot of theory" had to be in place before optogenetics could become a viable tool to develop, I point out that every piece of this necessary background theory itself resulted directly

from the development of a previous new tool for conducting experiments. In the actual history of laboratory sciences like contemporary neurobiology, the driving force behind everything is new tool development, “all the way down” (or back). Far from theory being primary, and so in desperate need of immediate development, in neuroscience new experiment tools have always come first. Theory tags along behind, completely at the behest of new experiment tools and their ingenious uses.

Attacking entrenched views like theory-centrism in (neuro-)science carries a strong burden-of-proof demand. Hacking, with his exclusive reliance on the single microscope example, and his cursory metascientific treatment of that case, surely failed to meet that burden. My strategy has been to investigate other case studies of the development of revolutionary experiment tools in neurobiology, to see if his conclusions fit those cases. So far, Hacking’s conclusions appear to fit with two historically revolutionary neuroscience tools, the metal microelectrode and the patch clamp (Bickle in preparation), as well as with my basic metascientific model derived from the gene targeting and optogenetics/chemogenetics cases.

I now want to carry this attack on theory-centrism in neuroscience one step further, by showing how such tools also illuminate Hacking’s claim, quoted above, that theory has only a “modest amount” to do with their development and that “it is engineering that counts.” If successful, that argument will really “put theory in its place,” as tertiary in importance in neuroscience. Theory will then have been shown to be completely dependent on the development and use of new experiment tools, and tool development completely dependent on engineering ingenuity. The principal case in my (2016), the development of gene targeting techniques, nicely illustrates Hacking’s further point. That argument will be my concern in Sect. 4 below. But gene targeting techniques were also a principal tool that brought about the mind-to-molecular-dynamics linkages that Kandel et al. alluded to in the quote that started this paper. So in the next section I’ll show how this tool has been used in neurobiology to forge these linkages “in a single bound.”

3 Gene Targeting Techniques in Molecular and Cellular Cognition (MCC)

Manipulating specific genes to “knock out” or “knock in” specific proteins in neurons *in vivo*, and tracing the behavioral effects, began in Seymour Benzer’s lab at Cal Tech in the 1960s. Only here the organisms were fruit flies, *Drosophila melanogaster*, and the behaviors tracked were mostly forms of simple olfactory conditioning. (Yes, fruit flies do learn to condition olfactory cues with electric shocks to resting surfaces.) By the late 1980s Benzer’s and other labs had developed more than 20 different learning and memory fly mutants, and had given them expressive names like ‘rutabaga’ and ‘dunce.’ Could this technique be extended to mammals, with their far richer behavioral repertoires? The possibility seemed daunting, due to the relative complexity linking gene expression and protein synthesis in mammals compared to insects. Then developmental biologists in the 1980s succeeded

in “knocking out” developmental genes in a mouse model using recombinant techniques. Two of those developers confidently asserted that those techniques should be applicable to any cloned gene (Thomas and Capecchi (1987)).¹

By the late 1980s a form of activity-dependent synaptic plasticity, one which strengthens or “potentiates” specific synapses between active neurons, provided a popular mechanism for some forms of learning and memory. The authors of the first systematic physiological study of this phenomenon speculated explicitly about its possible role in memory (Bliss and Lømo 1973), in part because the brain region from which they took tissue slices for their experiments, the mammalian hippocampus, had already been implicated in human learning and memory deficits. A decade-and-a-half after that initial study, such “long-term potentiation” (LTP) had been recorded in both tissue slices and in behaving rodent models for days up to weeks. The circumstantial case for this LTP-memory linkage was well known among neuroscientists (Lynch 1986). Post-synaptic N-methyl-D-aspartate receptors (NMDRs), to which the excitatory peptide neurotransmitters glutamate and glycine bind, provided a mechanism for some of the memory-like features of LTP. High activity in the pre-synaptic neuron was required for glutamate and glycine release; high activity in the post-synaptic neuron was required to remove a magnesium ion that blocked the ion channel in the NMDAR at or near resting potential. NMDARs thereby served as a kind of coincident-activity detector that neuropsychologist Donald Hebb (1949) had speculate about more than 3 decade prior. Hebb had envisioned mutually activated neurons forming temporary circuitries to realize various psychological functions. More importantly, activated NMDARs permitted the influx of not only sodium ions (Na^+), but also calcium ions (Ca^{++}) into the post-synaptic neuron. Few details were known at that time about what those Ca^{++} ions were doing, but their role in inducing LTP in post-synaptic neurons had already been established. Lynch et al. (1983) injected the calcium chelator ethylene glycol tetracetic acid (EGTA) into post-synaptic neurons under conditions normally inducing LTP. EGTA binds free Ca^{++} ions into inactive complexes. Its presence blocked LTP induction.

Still, the empirical case specifically for the LTP-memory linkage remained mad-deningly circumstantial. Morris (1989), whose name is indelibly attached to the rodent water maze task he perfected, provided the next key evidence. D,L-2-amino-5-phosphonopentanoic acid (AP5, also known as APV) is a potent and selective NMDAR antagonist. It blocks the influx of Ca^{++} ions into the post-synaptic neuron. Morris (1989) administered AP5 intraventricularly to rats prior to their learning a variety of hippocampus- and non-hippocampus dependent memory tasks (mostly using the Morris water maze). In rodent hippocampus tissue slices, AP5 was known to block LTP. Morris showed that it also decreased memory performance on the hidden platform version of his water maze, but not on the visual platform version. On the hidden platform task, hydrophobic rodents learn to find a platform submerged beneath the surface of a pool of opaque water by learning its location relative to

¹ I will say more about these techniques and their early development and successes in Sect. 4 below. The discussion to follow in this section draws on a more detailed account of gene targeting experiments in neurobiology in Bickle (2016, 3–9).

distal visual cues. This task is hippocampus-dependent; rodents with bilateral hippocampus lesions are impaired in the number of trials it takes them, compared to sham-lesioned controls. They also spend far less time in the quadrant of the platform's location on "probe" trials after acquisition, in which the platform is removed. On the visual platform task, a single visual cue is present at the location of the submerged platform (such as a flag sticking out the liquid). This task is not hippocampus-dependent; although rodents with bilateral hippocampus lesions are slower to learn it initially, by the end of standard training periods their performances match controls'. These were exactly the patterns of results Morris (1989) obtained in non-lesioned intraventricular AP5-treated rodents compared to vehicle-treated controls. Morris's Experiment Five also showed that doses of AP5 used in the water maze experiments were sufficient to block hippocampus LTP in vivo without interrupting normal synaptic transmission.

"The hypothesis that the physical substrate of memory in the mammalian brain resides in alterations of synaptic efficacy has been proposed frequently," Morris noted, "and is widely accepted by neuroscientists" (1989, 2052). With these new results, he suggested "that the type of synaptic plasticity studied in LTP experiments (1) is involved in some but not all kinds of learning, and (2) is involved in the initial associative phase of learning but not in retrieval" (1989, 2052). Notice that this is direct mind-to-cell-physiological-process linkage! Nevertheless, Morris himself noted that even highly selective pharmacological NMDAR antagonists like AP5 inevitably disrupt synaptic function in subtle ways, potentially interfering with activity throughout hippocampus circuitry. As Alcino Silva and collaborators put this worry 3 years later, perhaps the failure of learning Morris had painstakingly demonstrated "results not from the deficit in LTP, but simply from some other incorrect operation of hippocampus circuits that lack NMDA receptor function" (1992a, 201). Obviously, a selective NMDAR antagonist like AP5 could never unravel that potential confound.

Silva was aware of the predicted general applicability of Capecchi and colleagues' gene targeting techniques. He sought to use them to block LTP without disrupting other aspects of synaptic function, in order to better test the LTP-memory linkage. This approach would be novel. By "knocking out" the gene coding for some judiciously chosen protein product, an experimenter could abolish that protein's specific contribution to the phenomena of interest—LTP, rodent spatial learning and memory, whatever. The vast increase in specificity, compared to the best existing drugs, would hopefully not disrupt other aspects of synaptic function. But an immediate question loomed. Which neuronal genes/proteins should experimenters target, to test the LTP-memory linkage? And beyond that question were others. Was Thomas and Capecchi's (1987) speculation correct? Could gene targeting work for any cloned gene of interest in any type of tissue? In particular, could it work for a gene in neurons in vivo? Neurons are relatively delicate cells, susceptible to cell death in a variety of ways. The gene target chosen had to code for a protein significant enough in neuronal signaling pathways to block LTP if eliminated, and have downstream consequences all the way to behavior. These effects had to follow from the elimination of a *single gene* and its subsequent protein product. Yet this disruption also had to be specific enough not to disrupt other aspects of synaptic function

in excitatory forebrain neurons. Experimenters would have to verify that the targeted gene's transcription and protein production truly had been eliminated. The targeted gene could not interfere with normal development, since then-existing gene targeting techniques engineered the mutation at the embryonic stem cell stage, but tracked the behavioral effects in adult mice. Brain development in the mutants would have to be normal, from functioning excitatory synapses all the way up to gross anatomy of hippocampus circuitry and beyond. The behaving mutants would have to possess normal vision, motor capacities, and motivation to solve the behavioral tasks, so the engineered gene mutation would have to leave all those mechanisms intact. These were tall experimental demands!

Silva, in Susumu Tonegawa's lab, used these gene targeting technique to knock out the gene for the α -isoform of calcium-calmodulin-dependent kinase II (*α -CaMKII*). His choice was not random, of course. This protein was known to be highly enriched in the post-synaptic densities of mammalian forebrain excitatory neurons, including hippocampus and neocortex. It was known to play a role in NMDAR-dependent LTP. It is activated by calmodulin loaded with intracellular Ca^{++} , whose influx into the post-synaptic cell was through activated NMDARs. Activated α -CaMKII phosphorylates numerous other proteins known to be components of the mechanism of membrane depolarization, and itself remains activated via autophosphorylation after Ca^{++} influx ceases. In these ways α -CaMKII was already known to meet some key conditions on a computational model of a molecular mechanism to strengthen synapses (Lisman 1985; Lisman and Goldring 1988).²

So Silva et al. (1992a) constructed the plasmid, the small DNA molecule that replicates independently, to disrupt the *α -CaMKII* DNA sequence. They transfected the plasmid into mice embryonic stem cells, injected the stem cells into blastocytes, inserted the blastocytes into pseudo-pregnant females, and bred the resulting chimeric males with wild-type females. After multiple crosses confirmed the expected Mendelian ratios (wild-type homozygous, wild type-mutation heterozygous, mutation-homozygous) for what turned out to be a non-lethal engineered mutation, they showed that homozygous α -CaMKII mutants completely lacked α -CaMKII messenger RNA (mRNA) and protein in forebrain tissue. But these mice showed normal mRNA and forebrain protein levels for the closely related β -CaMKII isoform. Coronal sections through hippocampus revealed no gross anatomical abnormalities in cell types, distributions, or axonal pathways. Aside from "increased jumpiness" or "nervousness" when handled by humans, the mutant mice behaved normally. The mutation had no effect on long-term survival under standard laboratory housing. In light of all these preserved features, homozygous-wild-type littermates could be used as controls for experimental homozygous-mutant mice, for both slice-physiological and behavioral studies. The only difference between the mutant and wild-type mice was the absence or presence (respectively) of one protein, α -CaMKII, in neurons in forebrain regions.

² Lisman's model was a neurobiologically more specific version of Hebb's (1949) famous "neurons that fire together, wire together" speculation.

Electrophysiological studies using hippocampus slices found no differences between α -CaMKII mutants and controls in synaptic currents, for both NMDA and non-NMDA components. There were no differences in the dependence of NMDAR channel conductance on neuronal membrane voltage potentials (Silva et al. 1992a, Figs. 3 and 4). NMDAR function in α -CaMKII mutant mice slices was normal. So the targeted mutation did not disrupt synaptic function in any typically measured fashion, resolving the key confound plaguing the most careful pharmacological studies. But could this single targeted gene mutation produce the needed physiological and behavioral effects to directly link these molecular dynamics to memory?

Silva et al. (1992a) next investigated LTP in hippocampus tissue slices, using both field potential recordings to survey populations of hippocampus neurons and more sensitive whole-cell recordings. They demonstrated deficient LTP in mutant hippocampus slices, while littermate control slices showed normal tetanus-driven LTP for all time periods measured (up to 1 h after tetanizing stimulus to induce LTP). Aside from a brief post-tetanus stimulus potentiation (for about 1 min), synaptic strength in mutant slices was unchanged from baseline levels, and remained so even after a second tetanus with increased pulse trains was delivered. In whole-cell recordings most all hippocampus neurons in control slices exhibited normal LTP, while only a small fraction of neurons from α -CaMKII mutants did (Silva et al. 1992a, Figures 6, 7, 8, and Table 1). So this single gene/protein mutation reliably diminished LTP in hippocampus neurons.

Silva et al. (1992b) then investigated behavior in vivo in rodent hippocampus- and non-hippocampus-dependent learning and memory tasks. α -CaMKII mutants were slower to learn the non-hippocampus-dependent visible platform version of the Morris water maze task initially, but over a standard 2-day, 12-trial training period quickly matched littermate control performance. Interestingly, this pattern mimicked both hippocampus-lesioned and Morris 's (1989) AP5-treated animals (discussed above). On the hippocampus-dependent hidden-platform version, α -CaMKII mutants never learned to locate and mount the platform as quickly as controls, over either standard 3-day or 5-day training periods. In probe trials after training, where the platform is removed, mutants spent significantly less time in the maze quadrant where the platform had been located during training, and crossed the platform's training location significantly fewer times, than did controls. They also crossed the training location of the maze significantly fewer times compared to controls in a small number of random-platform trials interspersed during training, where the hidden platform was changed to a new location (Silva et al. 1992b, Figures 1, 2, 3, and 4). However, α -CaMKII mutants performed normally compared to controls on a water-filled plus maze task, which requires animals to use a single distal visual cue to learn which arm holds the hidden platform, (in contrast with the complex spatial relations between platform location and numerous distal visual cues, required on the hidden-platform Morris water maze task) (Silva et al. 1992b, Figure 5). Thus the mutants' failures to learn the hidden-platform Morris water maze task were not due to an inability to see the distal cues, or to learn an association between escape and the distal visual environment. Interestingly, in addition to their "jumpiness" upon human contact, α -CaMKII mutants also demonstrated other subtle behaviors similar

to mice with hippocampus lesions, such as increased exploration and activity in open fields and enclosed Y-mazes.

Silva's and collaborators' own words about their results, and about the potential of this gene targeting tool for neuroscience, are revealing. Their results "strengthen considerably the contention that the synaptic changes exhibited in LTP are the basis for spatial memory" (Silva et al. 1992b, 210). More surprisingly, and for the first time, they "demonstrate that a mutation in a known gene is linked to a specific mammalian learning deficit, and indicate that single genetic changes can have a selective but drastic impact on learning and memory" (1992b, 210). These, of course, are the mind-to-molecular-dynamics linkages that the quote from Kandel et al. at the beginning of this paper asserted. This experiment tool had now been used to test one such linkage, the role of α -CaMKII and the pathways it is part of to generate LTP in post-synaptic synapses, directly with hippocampus-based learning and memory tasks. Finally, and with an eye to the future use of this tool in mammalian behavioral neuroscience, they predict that "other similarly constructed mice with mutations in judiciously chosen genes will be useful for studying mammalian behavior" (Silva et al. 1992b, 210).

Neurobiologists did not have to wait long for this prediction to be met. Five months later a research team in Eric Kandel's lab, led by Seth Grant, published results from mice with engineered mutations to the genes for each of four nonreceptor tyrosine kinases (Grant et al. 1992). Grant and collaborators used the same gene targeting techniques to knock out the gene for each of these tyrosine kinases in different mouse mutants. Both electrophysiological and behavioral results with the *fyn* mutants exactly matched the results that Silva et al. (1992a, b) had achieved with their α -CaMKII mutants. (However, *fyn* mutants displayed a developmental neurological deficit, in the arrangements of hippocampus dentate gyrus granule neurons and their target CA3 pyramidal neurons; see Grant et al. 1992, Figure 7.) These authors were likewise enthusiastic about their results, and about the general applicability of this new gene targeting tool for neuroscience, which had now been shown to be feasible for mammalian behavioral neuroscience in two labs. "In addition to their role in the study of behavior and learning, targeted disruption of genes provides a powerful tool for examining the role of specific proteins in the function of the brain" (Grant et al. 1992, 1908).

In October 1992, 4 months after the Silva et al. papers had been published and 2 months before the Grant et al. paper appeared, Morris himself, with Mary Kennedy, published an invited review, "The Pierian Spring" in *Current Biology*, with the subtitle "mutant mice engineered to lack an enzyme critical for long-term synaptic plasticity are deficient in spatial learning" (Morris and Kennedy 1992, 511). They described for nonspecialists the genetic engineering procedures, and the electrophysiological and behavioral results achieved. They closed with a section on "Implications and potential." The α -CaMKII knock-out was "an ingenious piece of molecular engineering"; the LTP and behavioral deficits that matched hippocampus-lesioned and AP5-treated rodents "were by no means a foregone conclusion"; the work "should be recognized as the considerable achievement it truly represents"; and the key findings "vindicate and extend earlier results" (namely, Morris's own!) (1992, 513). Morris and Kennedy note that this approach was not without its own

problems. The experimenters might have underestimated the hippocampus-dependent learning capacities in the mutant mice. The subtlety of the effects of eliminating so significant a post-synaptic protein in forebrain excitatory neurons might reflect compensatory effects of other CaMKII isoforms. α -CaMKII had been eliminated completely, “knocked out,” from the mutants’ brains, including in pre-synaptic neurons, so these specific mutants could not be used to resolve the then-still raging controversy over whether LTP was mediated pre-synaptically, post-synaptically, or both. α -CaMKII is also prominent in neocortical excitatory neurons, and both consolidation and long-term storage of spatial memories probably occurs there, in addition to the role of hippocampus. So that phase of memory induction might also depend on α -CaMKII activation, which would also be blocked in the mutants. Lots of experimental questions still remained unanswered. Despite these issues, however, this first use of gene targeting techniques in behavioral neuroscience, according to Morris and Kennedy, was “an auspicious beginning and likely to fund a small industrial revolution” (1992, 514). Rather than advising neuroscientists to tread lightly with this new experiment tool, one of the world’s foremost behavioral neuroscientists advised his fellow experimentalists to “as Pope went on to write, “Drink deep, or taste not the Pierian Spring”” (1992, 514).³

Behavioral neuroscientists, and not just those working on learning and memory, certainly heeded Morris and Kennedy’s advice. Gene targeting quickly became standard practice in behavioral neuroscience, and the then-nascent search for molecular mechanisms of “higher” functions expanded quickly. This new experiment tool also helped bring molecular neuroscience to the prominence across the entire discipline it maintains to the present day. It provided a number of the early results that Kandel et al. were alluding to in the passage that begins this essay, Refinements of this basic tool also came quickly. Techniques were perfected for increased regional specificity of the targeted gene and its protein products, and for temporal specificity, the capacity to express engineered gene mutations only during specific developmental phases, e.g., in adult rodents, and even at specific times during behavioral tests. And a new molecular target for learning and memory studies also quickly emerged.

Günther Schütz’s developmental biology lab engineered a homozygous knock-out mouse, deleting the gene for the α - and δ -isoforms of cAMP-response element-binding protein (CREB). CREB is a transcriptional enhancer prominent in many biological tissues. Earlier work with flies, including gene targeting work, and with the sea slug *Aplysia californica*, suggested that CREB, when activated via phosphorylation, enhances expression and synthesis of a variety of genes and proteins important for LTP and for invertebrate forms of associative learning. Working with Schütz’s CREB mutant mice in Silva’s lab at Cold Spring Harbor, Roussoudan Bourtschuladze et al. (1994) showed they were deficient in long-term memory tasks (24-h delay), but intact in learning and short-term memory (30–60 min delay) on these same tasks.

³ Morris and Kennedy drew their title from this line of Alexander Pope’s early-eighteenth century poem, “An Essay on Criticism.” The line of the couplet that precedes it, also quoted by Morris and Kennedy at the beginning of the review, is perhaps the most famous line from the poem: “A little learning is a dangerous thing.”

Experimenters tests both hippocampus- and non-hippocampus-dependent tasks. These results suggested that CREB plays a key role in the consolidation of memory from short-term to long-term form. Hippocampus slice-physiology work proved consistent with this interpretation of the behavioral results. LTP in mutant slices was smaller than littermate controls and declined to baseline by 90 min (Bourtchuladze et al. 1994).

CREB mutant mice immediately became the target of extensive behavioral neuroscience investigations, and the same pattern of results emerged on a variety of rodent memory tasks: intact short-term performance but impaired long-term performance on the same tasks. This pattern held both for standard memory consolidation and for reconsolidation after stimulus re-presentation. The importance that *CREB* mutants played for learning and memory research over the next few years after Bourtchuladze et al.'s (1994) publication is nicely summarized by Yadin Dudai in his introductory sourcebook for memory research:

CREB is one of the most commonly used acronyms in neurobiology these days, and also one of the few words in the jargon of molecular biology that even experimental psychologists and computational neuroscientists might have encountered. And if they didn't, they should. Because the more we advance our knowledge in molecular biology the more we realize that CREB plays a pivotal role in the response of neurons to external stimuli. (2002, 65)

CREB is just one component of an intracellular signaling pathway that starts with neurotransmitter binding to receptors and ends with new gene expression, new protein synthesis, and reconstructed cytoskeletons at synapses. Dopamine, released by modulatory neurons, binds to G-protein-coupled receptors. The activated G-protein activates adenylyl cyclase to convert adenosine triphosphate into cyclic adenosine monophosphate (cAMP), a prominent second messenger in molecular biology. cAMP binds to regulatory subunits of protein kinase A (PKA) molecules. This binding frees catalytic PKA subunits to translocate to the neuron's nucleus and phosphorylate CREB molecules. Phosphorylated CREB in turn drives gene expression and synthesis of both regulatory and effector proteins which restructure active synapses, potentiating them to increased excitatory post-synaptic potentials to subsequent pre-synaptic glutamate release.⁴

Neuroscientists were also soon attracted to a different kind of gene targeting technique, the transgenic (or "knock-in") approach. This approach involves inserting an extra copy or copies of a cloned gene into the DNA of mammal embryonic stem cells, with the gene typically attached to a promoter region which limits its expression to specific neurons. Every cell in the mutants' bodies thus contains the extra transgene copy or copies; but transgene transcription (and subsequent protein synthesis) only occurs in those specific cells possessing the promoter molecule in sufficient amounts to turn on messenger RNA transcription. An early influential use of this approach in

⁴ This cell signaling pathway and its role in synaptic plasticity is so prominent in contemporary neuroscience that one finds detailed treatment of it in any good up-to-date textbook. I recommend Purves et al. (2018, chapter 7), especially for readers less familiar with molecular biology.

behavioral neuroscience targeting the cAMP–PKA–CREB–new gene expression–new protein synthesis pathway was Abel et al. (1997), in Kandel’s lab. They inserted extra copies of the gene for regulatory subunits of the cAMP-dependent PKA molecule into mouse embryonic stem cells, attached to a CaMKII promoter which limited transgene expression to forebrain regions, including hippocampus (but excluding significant expression in amygdala). In Abel et al.’s (1997) R transgene mutants, the extra R PKA (regulatory) subunits available in neurons in which the transgene is expressed quickly bind to PKA catalytic subunits freed by increased cAMP, blocking that early step in the cAMP–PKA–CREB pathway. As we saw above, CREB’s role in LTP and memory consolidation had already been established. But CREB is phosphorylated through numerous cell signaling pathways, and experiments with the CREB knock-out mice couldn’t distinguish between which of these pathways was crucial for its role in late-phase LTP and memory consolidation. Attempts to knock out the PKA gene using standard gene targeting techniques had been inconclusive.

Abel et al.’s founder mutant mice bred successfully and transmitted the transgene to offspring. The promoter limited significant transgene expression to forebrain areas, including all regions of the hippocampus, but hippocampus gross anatomy was otherwise unaffected. Hippocampus PKA activity was reduced in R transgenic mice, but synaptic transmission and early-phase LTP were unaffected. However, late-phase (L-) LTP, which requires new gene expression via CREB activation, was reduced significantly in mutant hippocampus slices. Behaviorally, R transgenic mutants were deficient compared to littermate control performance on both time-in-target-quadrant and number of target crosses in probe trails administered after training on the Morris water maze task (Abel et al. 1997, Figure 5). More importantly, R transgenic mutants were intact on short-term (1-h delay) contextual fear conditioning but significantly impaired on the long-term (24-h delay) version. But they were unimpaired on both short-term and long-term versions of tone-foot shock (Pavlovian) conditioning. The former task is hippocampus-dependent, where the R transgene was expressed significantly; the latter task is amygdala-dependent, where the R transgene was not expressed significantly.

These are just a handful of the earliest studies using just one of the novel tools neurobiologists have had available to link mind to molecular dynamics, in single bounds. This tool has undergone numerous refinements since these early studies. (See Silva et al. 2014 for a survey of landmark learning and memory results from the first 2 decades of the use of this tool in behavioral neuroscience.) And other tools for intervening cellularly/molecularly and tracking behavioral effects have been equally successful. But even this brief survey illustrates that the kinds of claims by neurobiologists, like the one from Kandel et al. with which this paper starts, are neither metaphorical nor hubristic. Neurobiologists have been linking mind to molecular dynamics for 3 decades now, using these new experiment tools.

4 Putting Theory in Its Place

I suggested above that the case study of gene targeting techniques shows more than just how neurobiologists have been linking mind to molecular dynamics for more than a quarter-century. It also shows the tertiary, not primary, role of theory in neurobiological practice. Linking behaviors indicative of cognitive functions directly to molecular mechanisms constitutes contemporary neuroscience's greatest, and best confirmed theoretical achievements. But these linkages are totally dependent on new experiment tools, like gene targeting techniques, which enabled experimenters to manipulate components of these molecular pathways in active, behaving animals. And now I want to point out how these experiment tools themselves came about. It wasn't from applying theory. It was from solving engineering problems. Hacking's old adage about microscopes, quoted above, that theory had only a "modest amount" to do with their development, and that it was "engineering that counted," applies very well to the tools that have revolutionized recent neurobiology and behavioral neuroscience.

Return to the initial development of gene targeting by homologous recombination in embryonic stem cells. This technique was applied to mammals by developmental biologists in the 1980s; the three most recognized scientists were Mario Capecchi, Martin Evans, and Oliver Smithies. Each was awarded a one-third share of the 2007 Nobel Prize for Physiology or Medicine "for their discoveries of principles for introducing specific gene modification in mice by the use of embryonic stem cells" (https://www.nobelprize.org/nobel_prizes/medicine/laureates/2007/). To make this tool work, embryonic stem cells had to be extracted and kept viable in cell cultures. They had to be inserted into host cells to engage the latter's recombinant mechanisms. Resulting "chimeric" mice had to survive, develop, and breed to produce offspring with the desired genetic mutation—and no "off-target" mutations. Mutant offspring had to survive and develop. And above all, the precise location of engineered gene lesion or transgene insertion had to be verified. The ways that Capecchi, Evans, and Smithies solved these many problems to build the first workable gene targeting technique for mammals display many signs of an "engineering-first" attitude that Hacking (1983) noted for microscopes.

For example, Capecchi reports that his "entry into what was going to become the field of gene targeting" began in 1977. He was "experimenting" with the use of tiny glass needles to inject DNA directly into the nuclei of cultured living cells (2007, 155). Neurobiologist Larry Okun, in the lab next door, was doing intracellular recordings of membrane potentials. Okun shared his recording apparatus with Capecchi and the tiny pipette turned out "to be ideal for conversion into a "microsyringe" to allow pumping of defined quantities of macromolecules, including DNA, into mammalian cells in culture" (Capecchi 2007, 155; see also Capecchi's "schematic" of his contrived DNA injection apparatus, Figure 2, 156). His new pattern of DNA delivery immediately improved the efficiency of stable incorporation of functional copies of new genetic material into cells, compared with then-existing techniques, by three orders of magnitude. He went from a success rate of roughly one cell out of one million to roughly one cell out of one thousand (2007, 155–156).

Engineering ingenuity with a tool designed for other purposes, even from a different scientific field, led to this first innovation.

Hacking (1983) emphasized the misleading role that theory often plays in generating initially negative reactions to a new tool's promise. He points to early skepticism about electron microscopes: "It was a long shot, because people were convinced, *on theoretical grounds*, that the specimen would almost instantly be fried and then burnt out" (1983, 199; my emphasis). Capecchi similarly recalls reviewers' reactions to his grant proposal to NIH in 1980, to test "the feasibility" of his lab's initial attempt to entrain the then-known mechanisms of homologous genetic recombination in mammalian cells to engineer specific genetic mutations. The project was "emphatically discouraged by the reviewers on the grounds that there might be only a vanishing small probability that the newly introduced DNA would ever find its matching sequence within the host cell genome" (2007, 160). Capecchi's proposal was denied funding. Despite this rejection and seemingly well-grounded theoretical rationale, Capecchi and colleagues "took a big gamble" and persisted in their attempts to develop a cell-selection method using the neomycin-resistant gene, co-inserted and activated only in cells in which the correct coupling of the inserted target gene to its matching sequence in the host cells occurred. Only those cells confer resistance to neomycin (by expressing the co-inserted neomycin-resistant product as well), and so only those cells survive *in vitro* in a neomycin-infused culture. Four years later the group reported evidence that successful gene targeting via recombinant mechanism occurs in cultured mammalian cells using their cell-selection method (Folger et al. 1984). Capecchi reports submitting another grant application to further develop this tool to the same NIH study section that had rejected his earlier proposal due to its presumed unfeasibility on molecular-biological theoretical grounds. "The section's response was 'we are glad that you didn't follow our advice'" (2007, 161).

Evans contributed to gene targeting technology by successfully isolating and sustaining mouse embryonic stem cells in tissue cultures *in vitro*. His remembrances nicely project the slow trial-and-error process and the "if it works, use it" mantra of an "engineering-first" attitude. He reports that around 1980 he was puzzling over why his initial attempts to grow his target cells directly from either explanted mouse embryos or dissected embryo tissues were failing. One of a handful of possibilities he considered was that "there might be only very small numbers of founder cells available and that therefore success *in vitro* would depend upon the highest efficiency of cloning" (2007, 185). He went with that single possibility. By that time, he notes, he had "slowly" improved cloning efficiency in a variety of other cell types, and "using this as a test for optimizing the media and conditions [had] arrived at a mix known around the lab at that time as 'Marvin's Magic Medium' or MMM" (2007, 185). Not exactly the application of theory to solve a problem! Evans also reports that the feeder layer he was using to maximize cloning efficiency "was also optimized by the same test" (2007, 185). Namely, trial-and-error, and an attitude of "if it works, use it."

Smithies is even more emphatic and explicit about an “engineering-first” attitude driving his contributions to gene targeting. He titled his Nobel Prize address “Turning Pages,” in reference to his personal laboratory notebooks, “more than 130 since I first began,” all of which he kept, from his almost 60-year career as a “bench scientist” (2007, 209). The many doodles and sketches from his notebooks that he reprints as part of his Nobel address beautifully illustrate Hacking’s (1983) point about the “pre-theoretical role of invention and fooling around” (1983, 199).⁵ Smithies reports that when his lab changed the cell types they had been working with, following continual failures to achieve homologous recombination, they discovered that their new cells “grew in suspension, and could only be transformed by a newly devised procedure—electroporation” (2007, 217). In the words of its discoverers’, electroporation is a method for introducing cloned genes into mammalian cells, “simple, rapid, and applicable to many (perhaps all) cell types, including those that are refractory to traditional transfection procedures,” whereby a suspension of cells and cloned DNA is exposed to a high-voltage electric discharge (Potter et al. 1984). The catch was that no electroporation machines were commercially available at that time. So Smithies “spent the next few months designing and testing a homemade apparatus, which was constructed inside a baby bathtub from part of a plastic test tube rack and electric parts from the local Radio Shack store” (2007, 217). His final version of the apparatus “does not look impressive—but it worked, and was subsequently used for all the definitive experiments” (Smithies 2007, 217). Smithies design schematic and photograph from his lab notebook (reproduced as Figure 9 in his 2007, 218) definitely meet his self-description above of “homemade.”

In fact, Smithies’ engineering-first attitude drove a number of his subsequent discoveries. After Evans had successfully isolated mouse embryonic stem cells, and both Smithies’ and Capecchi’s labs had successfully knocked out the *Hprt* gene in these cells using a drug-selection procedure to isolate correctly targeted cells in culture (described above in my discussion of Capecchi’s contributions), Smithies set about to find a more general procedure for targeting genes that did not have a directly selectable product. The trick, he guessed, would be to have available a simplified recombinant fragment assay of the targeted genetic material. He reports that he had recently attended a talk reporting the recently discovered polymerase chain reaction (PCR), which looked promising to him for this purpose. PCR is a method for synthesizing many copies of specific DNA sequences. In the words of its developers, it consists of

repetitive cycles of denaturation, hybridization, and polymerase extension and seems not a little boring until the realization occurs that this procedure is catalyzing a doubling with each cycle in the amount of the fragment defined by the positions of the 5’ ends of the two primers on the template DNA, that this fragment is therefore increasing in concentration exponentially, and that

⁵ See Bickle (2018) for a similar assessment and reproduction of a published page from Karl Diesseroth’s laboratory notebook depicting the initial schematic cartoon for the inserted light source for activating the light-sensitive receptor proteins used in optogenetics, a more recent tool in cellular physiology and behavioral neuroscience.

the process can be continued for many cycles and is inherently very specific. (Mullis et al. 1986).

Smithies once again faced the difficulty that “no suitable apparatus was commercially available ... so Hyang-Suk Kim and I made our own PCR machine,” which 20 years later “I still use” (2007, 223). The schematic design and a photograph of this machine, from Smithies’ lab notebook (2007, 224, Figure 15), will be amusing to anyone who has worked with a commercial PCR device (which are now routinely used in freshman biology teaching labs). When he later began using gene targeting techniques to investigate genetic factors in hypertension, Smithies reports employing “one of my glider pilot students,” John Smith, to rig up a computerized blood pressure measuring apparatus for mice. Why Smith? “I chose him to make the new machine ... because he had told me about a computerized device that he had designed and built to detect the stones left in pitted cherries, which cause lost teeth in the eaters and lawsuits against the suppliers” (2007, 225; see a photo of the cobbled-together device as Figure 18, 226). Once again, engineering first; the application of theory this isn’t.

It is important to remind ourselves that the outcome of all of this catch-as-catch-can laboratory tinkering was one of contemporary biology’s most revolutionary and widely-used experiment tools. The adoption and uses of gene targeting by homologous recombination in embryonic stem cells very quickly spread beyond developmental biology. As mentioned, Silva et al. (2014) provide a comprehensive survey of its landmark uses in behavioral neuroscience—in a book titled, not coincidentally, *Engineering the Next Revolution in Neuroscience*. The molecular mechanisms of cognitive functions rank among contemporary neuroscience’s greatest theoretical achievements. And yet this theory is tertiary in dependence. It comes directly from the development and ingenious experimental use of some novel experiment tools, to intervene into specific molecular processes in behaving mammals. And those tools come from a catch-as-catch-can, make-it-work, engineering-first attitude of the sort famously alluded to by Hacking (1983), in his “microscope” argument for the relative independence of “the life of experiment” from theory.

So what should we make of the incessant call for “more theory” in neuroscience? The lessons from a metascience of tool development in neurobiology on how to get that are straightforward: Build new and better experiment tools! Without notable exception, the most well-established, empirically secure theory about brain function has depended entirely on the development and ingenious experimental use of such tools. The mechanisms of neuronal conductance and transmission, the field properties of individual neurons, the detailed anatomical circuitries connecting neurons within and across brain regions, and increasingly the cellular and molecular mechanisms of higher cognitive functions—each of these accomplishments can be tied specifically to the development and use of one or more experiment tools. Not to arm-chair reflection about “how the brain works.” Neuroscience in anything resembling its past and current form is unimaginable without these practices. Practicing laboratory neuroscientists seem to recognize this point. The BRAIN Initiative—Brain Research through Advancing Innovative Neurotechnologies—launched in 2013, is a multi-U.S.-federal-agency, public–private collaboration “aimed at revolutionizing

our understanding of the human brain ... by accelerating the development and application of innovative technologies” (<https://braininitiative.nih.gov/>). The direction of influence and dependence this initiative is pursuing should be familiar to readers of this paper: engineering solutions → new experiment tools → better theory. Some scientists, including neuroscientists, have urged a deeper appreciation of engineering’s contributions to science’s progress (Diesseroth and Schnitzer 2013; Narayanamurti and Odumosu 2016). Cognitive neuroscientists and philosophers should take heed.

This paper has examined a single research tool, prominent in neuroscience for a quarter-century, and the surprising lessons it offers cognitive neuroscientists, cognitive scientists, and philosophers about neurobiologists’ talk of mind-to-molecular-dynamic linkages and the theory-centrism that still pervades many “higher level” approaches to the mind-brain. The engineering-first attitude it finds in the development of gene targeting techniques extends my previous work toward articulating a metascience of tool development experiments in neurobiology (Bickle 2016) and exploring the ways that these tools vindicate ideas about the independence of “the life of experiment” from theory first championed by Hacking (Bickle 2018). But these conclusions call out desperately for more case studies of the development of tools that have driven neuroscience’s most recognized successes. Nothing less than an alternate picture of how science works, drawing on a science-in-practice perspective on laboratory neuroscience, seems to be developing.⁶

References

- Abel T, Nguyen PV, Barad M, Deuel TAS, Kandel ER, Bourchouladze R (1997) Genetic demonstration of a role for PKA in the late phase of LTP and in hippocampus-based long-term memory. *Cell* 88:615–626
- Bickle J (2003) Philosophy and neuroscience: a ruthlessly reductive account. Springer, Dordrecht
- Bickle J (2015) Marr and reductionism. *Top Cogn Sci (TopiCS)* 7(2):299–311
- Bickle J (2016) Revolutions in neuroscience: tool development. *Front Syst Neurosci* 10:24. <https://doi.org/10.3389/fnsys.2016.00024>
- Bickle J (2018) From microscopes to optogenetics: Ian Hacking vindicated. *Philos Sci* 85(5):1065–1077
- Bickle J (in preparation) Tool development in neuroscience: putting theory in its place
- Bickle J, Kostko A (2018) Connection experiments in neurobiology. *Synthese* 195:5271–5295
- Bliss TV, Lømo T (1973) Long-lasting potentiation of synaptic transmission in the anesthetized rabbit following stimulation of the perforant path. *J Physiol* 232(2):331–356
- Bourchouladze R, Frenguelli B, Blendy J, Cioffi D, Schütz G, Silva AJ (1994) Deficient long-term memory in mice with a targeted mutation of the cAMP-responsive element-binding protein. *Cell* 79(1):59–68
- Capecchi MR (2007) Gene targeting: 1977–present. Nobel Lecture, pp 155–172. https://www.nobelprize.org/nobel_prizes/medicine/laureates/2007/capecchi_lecture.pdf. Accessed 18 May 2019
- Churchland PS (1986) *Neurophilosophy*. MIT Press, Cambridge
- Churchland PS, Sejnowski TJ (1992) *The computational brain*. MIT Press, Cambridge
- Churchland PS, Sejnowski TJ (2016) Blending computational and experimental neuroscience. *Nat Rev Neurosci* 17(11):667–668
- Craver CF (2007) *Explaining the brain*. Oxford University Press, New York

⁶ In recent talks I have offered two additional case studies that illustrate all these points: the metal microelectrode and the patch clamp. Bickle (in preparation) elaborates both cases.

- Crick FHC (1979) Thinking about the brain. In: Scientific American (ed) The brain. Scientific American Books. W.H. Freeman and Co, San Francisco, pp 130–140
- Crick FHC, Koch C (2003) A framework for consciousness. *Nat Neurosci* 6:119–126
- Diesseroth K, Schnitzer MJ (2013) Engineering approaches to illuminating brain structure and function. *Neuron* 80(3):568–577
- Dudai Y (2002) Memory from A to Z. Oxford University Press, New York
- Evans MI (2007) Embryonic stem cells: the mouse source—vehicle for mammalian genetics and beyond. Nobel Lecture, pp 181–192. https://www.nobelprize.org/nobel_prizes/medicine/laureates/2007/evans_lecture.pdf. Accessed 18 May 2019
- Folger KR, Thomas KR, Capecchi MR (1984) Analysis of homologous recombination in cultured mammalian cells. In: Cold spring harbor symposia on quantitative biology, vol 49, pp 123–138
- Gold I, Roskies AL (2008) Philosophy of neuroscience. In: Ruse M (ed) Oxford handbook of philosophy of biology. Oxford University Press, New York, pp 349–380
- Grant SGN, O'Dell TJ, Karl KA, Stein PL, Soriano P, Kandel ER (1992) Impaired long-term potentiation, spatial learning, and hippocampal development in *fyn* mutant mice. *Science* 258:1903–1910
- Hacking I (1983) Representing and intervening: introductory topics in philosophy of natural science. Cambridge University Press, Cambridge
- Hebb D (1949) The organization of behavior. Wiley, New York
- Kandel ER, Schwartz JH, Jessell TM (2001) Principles of neural science, 4th edn. McGraw Hill, New York
- Kuhn TS (1962) The structure of scientific revolutions. University of Chicago Press, Chicago
- Lisman JE (1985) A mechanism for memory storage insensitive to molecular turnover: a bistable autophosphorylation kinase. *Proc Natl Acad Sci USA* 82:3055–3057
- Lisman JE, Goldring MA (1988) Feasibility of long-term storage of graded information by the Ca²⁺/calmodulin-dependent protein kinase molecules of the postsynaptic density. *Proc Natl Acad Sci USA* 85:5320–5324
- Lynch G (1986) Synapses, circuits, and the beginnings of memory. MIT Press, Cambridge
- Lynch G, Larson J, Kelso S, Barrioneuvo G, Schottler F (1983) Intracellular injections of EGTA block induction of hippocampal long-term potentiation. *Nature* 305:710–721
- Marr D (1982) Vision. W.H. Freeman and Co, San Francisco
- Morris RG (1989) Synaptic plasticity and learning: selective impairment of learning rats and blockade of long-term potentiation in vivo by the N-methyl-aspartate receptor antagonist AP5. *J Neurosci* 9(9):3040–3057
- Morris RGM, Kennedy MB (1992) The pierriean spring. *Curr Biol* 2(10):511–515
- Mullis K, Faloona F, Sharf S, Horn G, Erlich H (1986) Specific enzymatic amplification of DNA in vitro: the polymerase chain reaction. In Cold spring harbor symposia on quantitative biology, vol 51. Cold Spring Harbor Laboratory Press, pp 263–273
- Narayanamurti V, Odumosu T (2016) Cycles of invention and discovery: rethinking the endless frontier. Harvard University Press, Cambridge
- Potter H, Weir L, Leder P (1984) Enhancer-dependent expression of human kappa immunoglobulin genes introduced into mouse pre-B lymphocytes by electroporation. *Proc Natl Acad Sci USA* 81(22):7161–7165
- Purves D, Augustine GJ, Fitzpatrick D, Hall WC, LaMantia A-S, Mooney RD, Platt ML, White LE (2018) Neuroscience, 6th edn. Oxford University Press, New York
- Robins SK (2016) Optogenetics and the mechanisms of false memory. *Synthese* 193:1561–1583
- Robins SK (2018) Memory and optogenetic intervention: separating the engram from the ephory. *Philos Sci* 85(5):1078–1089
- Silva AJ, Paylor R, Wehner JM, Tonegawa S (1992a) Impaired spatial learning in α -calcium-calmodulin kinase II mutant mice. *Science* 257:206–211
- Silva AJ, Stevens CF, Tonegawa S, Wang Y (1992b) Deficient hippocampal long-term potentiation in α -calcium-calmodulin kinase II mutant mice. *Science* 257:201–206
- Silva AJ, Landreth A, Bickel J (2014) Engineering the next revolution in neuroscience. Oxford University Press, New York
- Smithies O (2007) Turning pages. Nobel lecture, pp 209–230. https://www.nobelprize.org/nobel_prizes/medicine/laureates/2007/smithies_lecture.pdf. Accessed 18 May 2019
- Sullivan JA (2009) The multiplicity of experimental protocols: a challenge to reduction and non-reductionist models of the unit of neuroscience. *Synthese* 167:511–539
- Sullivan JA (2010) Reconsidering 'spatial memory' and the Morris water maze. *Synthese* 177:261–283

Sullivan JA (2018) Optogenetics, pluralism, and progress. *Philos Sci* 85(5):1090–1101

Thomas KR, Capecchi MR (1987) Site-directed mutagenesis by gene targeting in mice embryo-derived stem cells. *Cell* 51(3):503–512

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.