

# Chapter 9

## Long-Term Potentiation: One Kind or Many?

Jacqueline A. Sullivan

**Abstract** Do neurobiologists aim to discover natural kinds? I address this question in this chapter via a critical analysis of classification practices operative across the 43-year history of research on long-term potentiation (LTP). I suggest that this 43-year history supports the idea that the structure of scientific practice surrounding LTP research has remained an obstacle to the discovery of natural kinds as philosophers of science have traditionally conceived them.

### 9.1 Introduction

A unique aspect of being trained by Peter Machamer is the importance he places on understanding scientific practices, providing descriptively accurate accounts of those practices and teasing out the interesting philosophical implications. He also encourages his students to gain hands-on experience in those areas of science of interest to them. My experience working in a neurobiology laboratory as a graduate student in combination with his mentoring have been invaluable to my thinking on a host of philosophical issues, including the one I take up in this chapter: the relationship between neurobiological kinds and scientific practice.

To provide some relevant background, in a paper in 2009, I described an observation I made while undertaking research in a neurobiological laboratory. I worked in the field of synaptic plasticity, the ability of synapses to undergo changes in response to patterns of electrical activity, which is thought to underlie learning and memory. I noticed that neuroscience affords investigators the freedom to produce forms of synaptic plasticity in a plurality of different ways, using different patterns of afferent stimulation. Given that different stimulation protocols could potentially recruit different mechanisms, it was an open question whether phenomena produced in different laboratories using different protocols were the same phenomenon or different phenomena.

---

J.A. Sullivan (✉)  
Department of Philosophy, Rotman Institute of Philosophy,  
University of Western Ontario, London, ON, Canada  
e-mail: [jsulli29@uwo.ca](mailto:jsulli29@uwo.ca)

© Springer International Publishing AG 2017  
M.P. Adams et al. (eds.), *Eppur si muove: Doing History and Philosophy of Science with Peter Machamer*, The Western Ontario Series in Philosophy of Science 81, DOI 10.1007/978-3-319-52768-0\_9

127

One of the examples I used in my paper to illustrate the “multiplicity of experimental protocols” was long-term potentiation (LTP), which is generally defined as an activity-dependent increase in the strength of a synapse. I restricted my focus to then contemporary investigations of the role of a single protein kinase cascade in LTP, the extracellular-signal regulated kinase (ERK), and provided evidence that in response to different LTP induction protocols (e.g., theta-burst, high-frequency stimulation), all which investigators claimed could be used to successfully induce LTP, the kinase responded differently. I used this as a basis to suggest that it was a live possibility that different investigators who used different LTP induction protocols to investigate the mechanisms of LTP were actually investigating different phenomena produced by different mechanisms rather than the same phenomenon produced by the same or different mechanisms. In other words, given the structure of experimental practice, it was unclear if LTP was one kind or many kinds.

Some questions that I did not ask in the 2009 paper, however, were: (1) Do LTP researchers themselves take different instances of LTP to be the same kind of phenomenon or different phenomena? (2) Is there consensus in the field about how to “lump” or “split” the phenomena? (3) Is the multiplicity of experimental protocols an obstacle to the discovery of kinds that track actual divisions in the causal structure of the world (i.e., so-called “natural kinds”)?

In this chapter, I engage in an historical analysis of LTP in order to provide at least some preliminary answers to these questions. I begin in Sect. 9.2 with an analysis of two strategies that have been put forward in the philosophical literature on mechanisms for revising scientific taxonomies, what I will refer to simply as (a) the “natural kinds strategy” (Craver 2002) and (b) the “conventional kinds” strategy (Craver 2009). In Sect. 9.3, I use these strategies in combination with some conceptual tools for thinking about LTP experiments in order to answer the aforementioned questions.

## 9.2 Two Competing Constraints on Scientific Taxonomies

It is widely accepted that a primary aim of neuroscience is to describe the mechanisms that produce phenomena of interest (e.g., Bechtel 2008; Machamer et al. 2000; Craver 2007). Implicit in this account of mechanistic explanation is a characterization of scientific progress in which kinds of phenomena and the mechanisms productive of those phenomena change over time. The basic idea is that sciences like neuroscience begin by characterizing phenomena and organizing them into groups on the basis of detectable surface features. As empirical inquiry yields data concerning the mechanisms that produce these phenomena, revisions are made with the aim of accommodating that taxonomy to the mechanistic structure of the world. This process is iterative; as more is found out about the mechanisms underlying phenomena of interest, further revisions to the taxonomy are possible (e.g., Bechtel 2008; Bechtel and Richardson 1993; Craver 2007, 2009). This view is consistent with the idea that sciences like neuroscience have realist aims; they aim to provide explanations of phenomena that reflect “how actually” those phenomena are

produced in the natural world. In doing so, they approximate towards scientific taxonomies that reflect real divisions of kinds that correspond to the mechanistic structure of the world (See Craver 2006).

Areas of science that seek to accommodate their taxonomies of kinds to the mechanistic structure of the world may be described as being engaged in a search for so-called “natural kinds” (Craver 2002, 2007, 2009). In an early paper, Carl Craver (2002), noted that cognitive neuroscientists seemed to uphold a sufficient criterion for natural kindhood that he dubbed the “No Dissociable Realization (NDR)” condition.<sup>1</sup> On this condition, “instances of a natural kind have one and only one realizer” and “if there are two distinct realizers for a putative instance of a kind, there are really two kinds, one for each realizer” (Craver 2002, 962). The example that Craver provides to illustrate how this condition operates in practice is memory research. Although neuroscientists originally believed that memory was a single kind of phenomenon, findings from lesion studies and research on subjects with selective brain damage revealed that declarative memory depends on different brain structures than procedural memory. In light of these findings, the category of memory was split into two groups of phenomena. Further findings about memory mechanisms have prompted further subdivisions in the taxonomy of memory (e.g., Kandel and Squire 2008; Sweatt 2009). These and other examples of taxonomic practices in action (e.g., Bechtel 2008) lend support to the idea that cognitive neuroscientists implement the NDR condition when revising the taxonomy of kinds of memory. Thus, in instances where empirical inquiry yields findings that indicate that the current taxonomy of neurobiological kinds does not correspond to neural architecture, the taxonomy is revised so as to ensure such correspondence.

Craver (2009), however, offers an alternative account of the kinds of criteria that inform “taxonomic revisions” in neuroscience that is also compatible with mechanistic accounts of neuroscientific explanation. Whereas the NDR condition is consistent with the idea that neuroscientists treat the world as “an objective arbiter among competing taxonomies of kinds” (e.g., of memory) and that their explanatory interests play no role in shaping the kinds they ultimately discover, if we consider the situation from another perspective, claims Craver, we become privy to the fact that investigators themselves, rather than the world, determine when two kinds of phenomena or two kinds of mechanisms are different or the same. Consider again the declarative and procedural memory case. Empirical findings across animal and human studies indicate that declarative memory depends on structures in the medial temporal lobe and procedural memory depends on the basal ganglia. However, the first thing to notice is that the kinds of findings supporting this division emanate from different kinds of brains—animal and human brains—that differ from each other anatomically and mechanistically. Further, within a given human population and within the context of even a single experiment no two hippocampi or basal ganglia are alike—they vary in terms of a number of structural and constitutive details (e.g., overall shape of brain area, cell number, numbers of synaptic connec-

---

<sup>1</sup>I am using this condition that Craver puts forward in his 2002 paper as a heuristic because I think it gets something right about how some neuroscientists conceive of double dissociation experiments and what can be accomplished by using them.

tions, intracellular and extracellular molecular concentrations). Variations in behavioral performance and gross brain activity (e.g., BOLD signal) across subjects are also common in experiments in cognitive neuroscience.

The aforementioned differences could all be taken, from the perspective of an investigator, to correspond to different joints of nature. However, from the perspective of scientists, to acknowledge such differences as relevant for building a scientific classification system would result in an unwieldy taxonomy of kinds that would contain as many different kinds as there are different individuals (either animals or humans). This would make the taxonomy wildly intractable and bar the kind of generality that neuroscientists desire for their explanations. Investigators must instead strike a delicate balance between “characterizing [a] mechanism very abstractly”, which “potentially glosses over sub-kinds of mechanism” and “characterizing [a] mechanism in maximal detail”, which “threatens to make each particular mechanism a kind unto itself” (Craver 2009, 587; See also Craver 2014).

Striking such a balance is, in fact, what cognitive neuroscientists may be characterized as doing. The taxonomic division between procedural and declarative memory abstracts away from certain differences while acknowledging others. The upshot is that rather than adhering strictly to developing taxonomies to directly mirror the causal structure of the world, “judgments about whether two mechanisms are mechanisms of the same kind rely ineliminably on judgments by people (in concert) about the appropriate degree of abstraction required for the problem at hand” (Craver 2014, 589). In other words, depending upon what the aims of inquiry are, be they explanation, prediction or intervention, the way investigators carve up the world will be aligned with their goals.

Even if we acknowledge, as Craver does, that conventional factors play a role in the development of our scientific taxonomies, this does not mean that investigators in neuroscience are not still aiming independently or collectively to discover something like or close enough (from their perspectives) to natural kinds. Craver seems sympathetic to this idea (See for example, Kendler et al. 2011). However, as I aim to show via an investigation of some relevant highlights of the 43-year history of LTP research, sometimes experimental practice is not conducive to the realization of this goal and investigators end up discovering kinds that are closer to the conventional kinds (antirealist) rather than natural kinds (realist) end of the kinds continuum (Craver 2009).

### 9.3 LTP: One Kind or Many?

Since at least the early twentieth century psychologists and physiologists hypothesized that associative forms of learning required changes in how neurons communicate with each other (See Kandel and Spencer 1968). In 1940, the psychologist Donald O. Hebb, proposed a mechanism for such changes. Hebb claimed that when two cells, A and B, which communicate across a synapse under normal conditions, undergo a period of repeated and concurrent activation, the result is a strengthening

of the connection between the two cells, exhibited by a subsequent change in the way cell A excites cell B (Hebb 1949). According to Hebb, each learning event is accompanied by a brief associated activation of two neurons that comprise a synapse, which results in the memory of that event being stored in the form of a physiological change at that synapse.

Although Hebb's proposed mechanism for learning, as one among several "cellular-connection theories of learning" (Kandel and Spencer 1968, 68), had many supporters, during the 1950–1960s, investigators tried to produce activity-dependent Hebbian-like changes at several nervous system synapses (e.g., spinal cord, lateral geniculate nucleus) with limited success.<sup>2</sup> In 1966, however, in the context of investigating the physiology of the dentate gyrus in the hippocampus of adult anesthetized rabbits, Terje Lømo observed an artificially induced physiological equivalent of a strengthening in synaptic efficacy much like the one Hebb had described (See Lømo 2003). After applying a brief yet repetitive stimulation of "one-second bursts of high-frequency (100 Hz "tetanic") stimulation" to perforant path fibers that project from the entorhinal cortex to the dentate gyrus, he recorded, extracellularly, an enduring observable increase in the amplitude above baseline of the evoked field potentials of post-synaptic dentate granule cells. This finding prompted him and his colleagues to further investigate the phenomenon (e.g., Bliss and Lømo 1973; Bliss and Gardner-Medwin 1973) and culminated in the publication of a now famous paper that introduced the scientific community to "a long-lasting potentiation of synaptic transmission" (Bliss and Lømo 1973)—the phenomenon later renamed "long-term potentiation" (LTP).

In order to produce LTP in the dentate area, Bliss and Lømo selected two locations in the hippocampus in which to place stimulating electrodes. One electrode was placed in the lateral perforant path and was used to deliver test pulses (0.1 msec, with maximum amplitude of 100 V) prior to LTP induction. The second electrode was placed in the medial perforant path and used to deliver LTP-inducing stimulation. They selected two locations from which to record extracellularly: (1) the cell body layer, where they recorded population spikes and (2) the dendritic layer, from which they recorded excitatory post-synaptic potentials. Once the stimulating and recording electrodes were at their desired locations in the brain, they applied test pulses consisting of "single shocks at a fixed strength, repeated at intervals of 2–3 sec[onds]" (Bliss and Lømo 1973, 334). Responses were recorded at regular intervals for up to 30 min, "with average responses based on 20 or 30 consecutive single responses" (Bliss and Lømo 1973, 334). A "sequence of conditioning trains" was then delivered at "intervals of 30 min or more". In each experiment they applied

---

<sup>2</sup>Little work had been done to study cortical synaptic plasticity in the mammalian brain due in part to technological limitations (Kandel and Spencer 1968, 85–86). However, quite a bit of work had been undertaken to induce changes in synaptic efficacy in the invertebrate, *Aplysia depilans* (Kandel and Spencer 1968). Early work on *Aplysia* indicated that activity-dependent changes in synaptic strength primarily involved pre-synaptic (e.g., changes in neurotransmitter release) as opposed to post-synaptic mechanisms (Kandel and Spencer 1968; See also Sweatt 2016).

“one or more conditioning trains” with the trains delivered at “10–20/sec[onds] for 10–15 sec[onds] or at 100/sec[ond] for 3–4 sec[onds]” (Bliss and Lømo 1973, 331).

Bliss and Lømo investigated several parameters of dentate granule cell responses, which they measured, following application of LTP-inducing conditioning trains to perforant path fibers: (1) changes in the amplitude of the population EPSP above average baseline response to test pulses; (2) changes in the peak of the amplitude of the population spike above baseline response to test pulses; and (3) latency of the population spike from post-conditioning test pulses to initial peak of the population spike. Across 15 experiments, they determined that “all three parameters were potentiated in 29% of the experiments” (1973, 331) they conducted. As evidence that they had detected a *bona fide* potentiation, they pointed to the fact that (1) the amplitude of the population EPSP increased in 43% of all experiments, (2) the population spike increased in 40% of all experiments and (3) a reduction in latency of the population spike occurred in 57% of all experiments. Their interpretation of the results was that “two independent mechanisms [were] responsible for [the] long-lasting potentiation” that they observed, namely: “(a) an increase in the efficiency of synaptic transmission at the perforant path synapses” and “(b) an increase in the excitability of the granule cell population” (1973, 332).

Immediately on the heels of Bliss and Lømo’s discovery, other investigators began producing LTP in their laboratories. The number of publications directed at understanding the phenomenon and its mechanisms exceeded 3000 by 1999 and 6000 by 2004 (See Sweatt 2016). Before considering some of the details of LTP research during this historical timeframe that is relevant to the question of what kind of kind LTP is, it is relevant to first identify some general features of LTP experiments and relate them to the concepts of natural and conventional kinds described in the previous section.

Experimentation consists of two stages (a) data production and (b) data interpretation (e.g., Woodward 2003). In LTP experiments, data production may be further subdivided into two stages: (1) a design stage in which an experiment is designed and a stimulation paradigm and subprotocol for producing LTP are selected and (2) an implementation stage in which the experimental design is carefully instantiated across each individual experiment. The immediate output of each individual experiment (or implementation of the experimental design) is an individual data point or set of data points. Investigators appeal to these data in order to adjudicate among competing hypotheses about LTP (e.g., hypotheses about its phenomenological features, its synaptic and cellular and molecular mechanisms) (See Sullivan 2009).

As was mentioned above, in the early stages of developing a scientific taxonomy, investigators group together kinds of phenomena that they take to be similar in terms of observable features. Taxonomies are subsequently revised in response to discoveries pertaining to the mechanisms productive of those phenomena. Thinking about the aforementioned features of LTP experiments provides additional insights into the processes that contribute to stabilizing scientific taxonomies and the phenomena to which they correspond. Bliss and Lømo selected an experimental design and sub-protocol for producing LTP. Insofar as across their individual experiments they adhered tightly to this experimental design, they took themselves to be

producing, measuring and detecting the same phenomenon across all of the experiments that had identical protocols they undertook in the laboratory. In doing so, they abstracted away from certain differences across these different experiments—different experimental subjects (e.g., rabbits, rats), different times of day in which the experiments were conducted, different stimulating and recording electrodes, different hippocampi, different perforant paths, different granule cells. Insofar as they regarded the data production processes across all of these experiments as reliable, they took themselves to be warranted in grouping the effects produced as instances of the same phenomenon.

Yet, were Bliss and Lømo warranted in grouping the effects produced across their different LTP experiments as instances of the same kind of phenomenon? I think our intuition is to respond “yes” to this question in part because we assume that just so long as the experimental process across these different experiments was reliable, then the mechanisms productive of the effects produced across them were sufficiently identical to be able to warrant classifying them as the same phenomenon (See Feest 2011 on the conditions under which phenomena are stabilized). When scientists make such judgments, they appear to be applying something like the NDR criterion; insofar as they think that the mechanisms across each experiment are identical or conserved, there are no grounds for splitting LTP into as many different kinds as there are LTP experiments. In other words, the scientists regard such abstractions away from the specific details of these experiments as legitimate just so long as they have good grounds for believing that the measures they have taken across experiments do serve to track or preserve real differences. Even though it is clear that conventional factors (i.e., decisions on the part of two investigators as to how to produce, detect and measure the phenomenon) are playing a role in shaping the kinds of phenomena discovered, there seems to be some sense in which the different instances grouped together constitute a *bona fide* kind of phenomenon from the perspective of the investigator(s).

What happens, however, when we consider LTP research taking place after the publication of Bliss and Lømo’s 1973 paper? Immediately on the heels of their discovery, other neurophysiologists began producing what they regarded as forms of mammalian LTP in their laboratories. Bliss and Gardner-Medwin, for example, went on to demonstrate that same year that “the same phenomenon” could be produced “without anaesthesia and under normal more stable conditions which can be obtained using chronically prepared animals” (1973, 358). Robert Douglas and Graham Goddard sought to “repeat the observations of Bliss and Gardner-Medwin using the rat instead of the rabbit” but made “several modifications to the procedure” in order to “make the observed potentiation more reliable” (1975, 206). By 1975, other neurophysiologists were producing LTP in their labs, using *in vitro* brain slice preparations (e.g., Deadwyler et al. 1975) and other novel electrophysiological recording techniques, including patch clamp (See Sweatt 2016). By 1976, *in vitro* experiments had been used to produce LTP in three subregions of the hippocampus (CA1, CA3 and dentate gyrus), and “the time course of appearance, magnitude of the effect and duration of LTP appear[ed] to be similar in all 3 areas” (Alger and Teyler 1976, 469). In 1976, Gary Lynch, V. Gribkoff and S. Deadwyler

published a letter to the journal *Nature* indicating that “hippocampal synapses” displayed “an unusual degree of physiological plasticity” and that the findings that “modest levels of repeated stimulation cause considerable enhancement in subsequent responses to single pulse stimulation” had “been replicated in several laboratories” (1976, 151).<sup>3</sup>

From the late 1970s to the mid-1980s, the LTP field exploded exponentially. By 1979, LTP at hippocampal synapses had become a model for understanding LTP in the nervous system more broadly and a lot of work was already being directed at uncovering the biochemical processes (e.g., second-messenger signaling cascades) and neurotransmitters (e.g., glutamate) involved in LTP at hippocampal synapses (e.g., Browning et al. 1979; Lynch et al. 1979; Dunwiddie and Lynch 1978; Baudry and Lynch 1980; Dolphin et al. 1982, 287). Different investigators used different stimulation paradigms and subprotocols, some that were intended to mimic “naturally occurring firing patterns, observed *in vivo* in the hippocampus” (e.g., theta-burst stimulation) and others that were used to induce LTP “by pairing repeated, single presynaptic stimuli with postsynaptic membrane depolarization” (“so-called ‘pairing’ LTP”) (Sweatt 2016).

I want to assume, for the sake of argument that investigators who engaged in this early LTP research regarded the instances of LTP they were producing in their own laboratories as *bona fide* instances of LTP, much like Bliss and Lomo did, and for similar reasons. That is, with respect to each published research study, just so long as investigators had ensured the reliability of their LTP-producing experiments, they believed themselves justified in thinking that they had stabilized a single kind of phenomenon in their own laboratories. However, if we compare and contrast different features of these experiments—different experimental subjects (rabbits, rats, guinea pigs), different stimulation paradigms and protocols for producing LTP, different preparations (*in vitro*, *in vivo*), different time course with respect to how long the potentiated effect lasted, it becomes difficult to assess whether all of these separate laboratory effects are instances of the same phenomenon or different phenomena. Are the mechanisms across experiments that use different model organisms, different hippocampal synapses or different stimulation paradigms and protocols or different preparations all the same?

In the 1970s and 1980s, investigators did not have answers to these questions.<sup>4</sup> Not enough was known about the cellular and molecular mechanisms of LTP, nor whether cellular and molecular activity differed depending upon the organisms being investigated or stimulus parameters and preparations being used. In response to such uncertainties, some investigators qualified the kind of LTP they were investigating by pointing to differences in the stimulation paradigms, animals and preparations used (*in vivo*, *in vitro*), and/or the synapses at which LTP-inducing

---

<sup>3</sup>The general claim that persisted in the literature until 1986 was that “long-term potentiation [. . .] seen in several hippocampal pathways following repetitive stimulation, [was] somewhat unique when compared to the post-tetanic potentiation seen at the neuromuscular junction or in invertebrates” (Dunwiddie and Lynch 1978, 353–354) in so far that it was longer lasting.

<sup>4</sup>They still lack answers, as I explain later in this section.



stimuli had been delivered.<sup>5</sup> Other investigators, at least prior to 1986, seemed amenable to the idea that despite these kinds of differences, the same phenomenon was under study at least across all hippocampal LTP experiments (e.g., Dunwiddie, Madison and Lynch 1978; Lynch et al. 1979). In other words, there was no real consensus as to how to “lump” or “split” the phenomena. Given that LTP research was, from 1973–1988, in what some reviewers later characterized as “a descriptive phase” (See for e.g., Nicoll, Kauer and Malenka 1988, 97)—a phase in which little was being learned about the mechanisms of LTP—this general lack of consensus about how to “lump” or “split” the phenomena makes sense (See also Craver 2003).

By 1988, LTP research had entered “a mechanistic phase” and results from a series of experiments that involved blockade of N-methyl-D-aspartate (NMDA) receptors at different LTP synapses were taken to establish that LTP at mossy fiber synapses did not require activation of NMDA-receptors (e.g. Harris and Cotman 1986; Staubli 1992).<sup>6</sup> This prompted investigators to split the category of hippocampal LTP into two broad subcategories: NMDA-receptor dependent and NMDA-receptor independent LTP (See also Nicoll et al. 1988; Nicoll and Malenka 1995). Notice that this suggests that there is some sense in which neurobiologists were upholding something like the NDR criterion—given the discovery that LTP could be produced with and without NMDA receptor activation, the hippocampal LTP category was split.

By 1994, investigators had begun to acknowledge that LTP could be induced in many different synaptic pathways by a variety of induction paradigms, and that the biochemical mechanisms of these forms of LTP may differ (See Malenka 1994; See also Lisman et al. 2003). Despite such admissions, some investigators still grouped results from different experiments together in order to make general claims about specific forms of LTP. For example, in 1999, Robert Malenka and Roger Nicoll found it necessary to restrict the focus of their review to LTP induced at “synapses between the Schaffer collateral and commissural axons and the apical dendrites of CA1 pyramidal cells” (1999, 1870). They also acknowledged that because “a review of the literature generates an enormous, even bewildering, list of candidate signal

---

<sup>5</sup> Bliss and Gardner-Medwin, for example, acknowledged, “the mechanisms of the effect remain uncertain” (Bliss and Gardner-Medwin 1973, 373). Douglas and Goddard were careful to indicate, “this type of potentiation may underlie memory storage in one part of the mammalian brain”<sup>5</sup> (1975, 214) rather than at all mammalian cortical synapses. While some investigators regarded “intracellular recordings from *in vitro* preparations of both immature and mature hippocampal tissue” as “similar to those obtained *in vivo*” (Deadwyler et al. 1975, 80), one early failure to obtain LTP in the dentate gyrus of the hippocampus was attributed to the possibility that *in vitro* slice preparations could compromise the integrity of the synaptic pathways (Deadwyler et al., 1975, 84) and result in a “decreased amount of recurring excitation” compared to *in vivo* preparations (Alger and Teyler 1976, 478). By 1978, Dunwiddie and Lynch determined that “various conditioning frequencies apparently induce[d] different degrees of long-term potentiation” (Dunwiddie and Lynch 1978, 366) and that synaptic transmission was required for the initiation of LTP (Dunwiddie, Madison & Lynch 413).

<sup>6</sup> Ursula Staubli, for example, claimed “mossy fiber potentiation is unlike LTP both in induction and expression mechanisms and thus is a wholly different form of synaptic plasticity” (Staubli 1992, 151).

transduction molecules” involved in LTP at these synapses, they were focusing only on those results concerning the mechanisms of this form of LTP that they took to be “compelling” (1999, 1871). By 2003, concerns about whether the mechanisms of different forms of NMDA-receptor dependent LTP (alone) differed depending upon the LTP stimulation paradigms and protocols used were prevalent in the LTP community. Specifically, Robert Malenka pointed to what was then becoming “the increasingly popular hypothesis that different LTP induction protocols result in mechanistically distinct types of NMDA receptor-dependent LTP”, which was “often used as a polite explanation for discrepancy in results between [research] groups”. Although Malenka suggested that “the evidence in support of this idea remain[ed] weak”; he acknowledged that “it will remain important to seriously consider the possibility that the patterns of activity that are used to elicit LTP influence which intracellular signalling cascades are activated” (Malenka 2003, 925). The following year, Malenka and Mark Bear indicated that “when discussing LTP [...] it is now necessary to define at which specific synapses these phenomena are being studied, at what time point during development, and how they are being triggered (Malenka and Bear 2004, 5). As partial indication that such issues have not yet been fully resolved, Baudry and colleagues recently claimed, “as has been repeatedly mentioned, a major difficulty to integrate all the findings” about LTP “is due in part to the use of different experimental protocols by the majority of research laboratories working on this topic, and the apparent lack of reproducibility of experimental data resulting from these differences” (2015, 74).

Given even these spotty details of the history of LTP, an interesting picture emerges with respect to the kinds of constraints that have to date informed the development of a taxonomy (or lack thereof) of LTP. Over the course of a 43-year history, the published record included well over 6000 papers on LTP. If investigators had begun to treat each independent instance of LTP produced in a given laboratory as an independent phenomenon having an independent mechanism, they would have risked having to develop an unwieldy taxonomy containing as many different kinds of LTP as there were experiments for producing it. That clearly was not practical. So, researchers instead (at least in the context of review papers) abstracted away from specific details of experimental practice that may have led to differences in mechanisms in order to make the goal of providing a unified model of LTP or unified models of specific forms of LTP (e.g. hippocampal NMDA-receptor LTP in area CA1 tractable). Phenomena that may indeed have been produced by different mechanisms thus were sometimes treated as instances of the same phenomenon.

However, notice that by treating different instances of LTP that may have had different mechanisms as all the same phenomenon, researchers abstracted away from potentially relevant causal differences, which seems antithetical to the discovery of kinds that track real divisions in the causal structure of the world.<sup>7</sup>

---

<sup>7</sup>At the time of Bliss and Lomo’s discovery, LTP had been reliably produced in invertebrates across many different laboratories. Bliss and Lomo’s finding was different, because it was produced in the mammalian brain. Debates about whether the cellular and mechanisms for LTP induction were conserved across species persisted well into the 21st century (See for example Bickel 2006). It was

That such pragmatic factors were operative in the development of the LTP taxonomy remained hidden until recently, when more and more investigators began to reach consensus that maybe the forms of even hippocampal area CA1 LTP being produced in different laboratories were really different as opposed to the same phenomenon. The reason that they remained hidden, in part, is that there were discoveries, like the discovery of non-NMDA-receptor-dependent LTP that seemed to confirm the idea that LTP researchers were in search of natural kinds and inclined to uphold the NDR criterion. However, as I have demonstrated in this section, a closer look at LTP research reveals that pragmatic factors have played a prominent role in shaping the current taxonomy of LTP.

## 9.4 Conclusion

At the start of this chapter, I raised four questions to which I want to return in light of the aforementioned analysis. First, do LTP researchers take different instances of LTP produced across different experimental protocols to be the same kind of phenomenon or different phenomena? I think there is widespread recognition that these different instances may not be the same phenomenon and that there is no real consensus about how to “lump” or “split” the phenomena. To cope with the vastness of the experimental record on LTP, some researchers have abstracted away from differences in experimental protocols in an attempt to provide unified mechanistic models of this or that specific form of LTP (e.g., Malenka and Nicoll 1999; Baudry et al. 2015). So, what kind of kinds are LTP researchers aiming to discover? From one perspective it may indeed be accurate to say that investigators are interested in discovering the mechanisms of LTP and developing a taxonomy that corresponds to real divisions in kinds of LTP. However, the multiplicity of experimental protocols, combined with the fact that discovering mechanisms is a collaborative enterprise are factors that have contributed to neurobiologists having to make pragmatic decisions as to how to lump or split the phenomena.

**Acknowledgement** The author would like to thank Uljana Feest for helpful comments on an earlier version of this paper and Floh Thiels and Peter Machamer for many interesting and helpful discussions about LTP.

---

common to hear some investigators arguing that LTP involved pre-synaptic mechanisms and others that it involved post-synaptic mechanisms. Support for conservation of mechanisms waxed and waned depending on the grain of analysis one used to assess similarities and differences in mechanisms across organisms (See Bechtel and Mundale 1999, Sullivan 2009; Craver 2009).

## Bibliography

- Alger, B., and T. Teyler. 1976. Long-Term and Short-Term Plasticity in the CA1, CA3 and Dentate Regions of the Rat Hippocampal Slice. *Brain Research* 110: 463–480.
- Anderson, P., and T. Lømo. 1966. Mode of Activation of Hippocampal Pyramidal Cells by Excitatory Synapses on Dendrites. *Experimental Brain Research* 2: 247–260.
- Baudry, M., and G. Lynch. 1980. Hypothesis Regarding the Cellular Mechanisms Responsible for Long-Term Synaptic Plasticity in the Hippocampus. *Experimental Neurology* 68(1): 202–204.
- Baudry, M., G. Zhu, Y. Liu, Y. Wang, V. Briz, and X. Bi. 2015. Multiple Cellular Cascades Participate in Long-Term Potentiation and in Hippocampus-Dependent Learning. *Brain Research* 1621: 73–81.
- Bechtel, W. 2008. *Mental Mechanisms: Philosophical Perspectives on Cognitive Neuroscience*. New York: Taylor and Francis.
- Bechtel, W., and J. Mundale. 1999. Multiple Realizability Revisited: Linking Cognitive and Neural states. *Philosophy of Science* 66(2): 175–207.
- Bechtel, W., and R. Richardson. 1993. *Discovering Complexity: Decomposition and Localization as Strategies in Scientific Research*. Princeton: Princeton University Press.
- Bennett, M. 2000. The Concept of Long-Term Potentiation of Transmission at Synapses. *Progress in Neurobiology* 60: 109–137.
- Bickle, J. 2006. Reducing Mind to Molecular Pathways: Explicating the Reductionism Implicit in Current Cellular and Molecular Neuroscience. *Synthese* 151: 411–434.
- Bliss, T., and R. Gardner-Medwin. 1973. Long-Lasting Potentiation of Synaptic Transmission in the Dentate Area of the Unanesthetized Rabbit Following Stimulation of the Perforant Path. *Journal of Physiology* 232: 357–374.
- Bliss, T., and T. Lømo. 1973. Long-Lasting Potentiation of Synaptic Transmission in the Dentate Area of the Anaesthetized Rabbit Following Stimulation of the Perforant Path. *Journal of Physiology* 232(2): 331–356.
- Bogen, J., and J. Woodward. 1988. Saving the Phenomena. *The Philosophical Review* 97: 303–352.
- Browning, M., T. Dunwiddie, W. Bennett, W. Gispen, and G. Lynch. 1979. Synaptic Phosphoproteins: Specific Changes After Repetitive Stimulation of the Hippocampal Slice. *Science* 203(4375): 60–62.
- Craver, C. 2002. Dissociable Realization and Kind Splitting. *Philosophy of Science* 71(5): 960–971.
- . 2003. The Making of a Memory Mechanism. *Journal of the History of Biology* 36: 153–195.
- . 2006. When Mechanistic Models Explain. *Synthese* 153: 355–376.
- . 2007. *Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience*. Oxford: Oxford University Press.
- . 2009. Mechanisms and Natural Kinds. *Philosophical Psychology* 22(5): 575–594.
- . 2014. The Ontic Account of Scientific Explanation. In *Explanation in the Special Sciences: The Case of Biology and History*, Synthese Library, Vol. 367. Dordrecht: Springer.
- Craver, C., and L. Darden. 2001. Discovering Mechanisms in Neurobiology: The Case of Spatial Memory. In *Theory and Method in the Neurosciences*, ed. P.K. Machamer, R. Grush, and P. McLaughlin. Pittsburgh: University of Pittsburgh Press.
- Deadwyler, S., F. Dudek, C. Cotman, and G. Lynch. 1975. Intracellular Responses of Rat Dentate Granule Cells *in vitro*: Posttetanic Potentiation to Perforant Path Stimulation. *Brain Research* 88: 80–85.
- Dolphin, A., M. Errington, and T. Bliss. 1982. Long-Term Potentiation of the Perforant Path *in vivo* is Associated with Increased Glutamate Release. *Nature* 297: 496–497.
- Douglas, R., and G. Goddard. 1975. Long-Term Potentiation of the Perforant Path-Granule Cell Synapse in the Rat Hippocampus. *Brain Research* 86: 205–215.