



Recharacterizing scientific phenomena

David Colaço¹

Received: 24 January 2019 / Accepted: 4 February 2020 / Published online: 25 February 2020
© Springer Nature B.V. 2020

Abstract

In this paper, I investigate how researchers evaluate their characterizations of scientific phenomena. Characterizing phenomena is an important – albeit often overlooked – aspect of scientific research, as phenomena are targets of explanation and theorization. As a result, there is a lacuna in the literature regarding how researchers determine whether their characterization of a target phenomenon is appropriate for their aims. This issue has become apparent for accounts of scientific explanation that take phenomena to be explananda. In particular, philosophers who endorse mechanistic explanation suggest that the discovery of the mechanisms that explain a phenomenon can lead to its recharacterization. However, they fail to make clear how these explanations provide warrant for recharacterizing their explananda phenomena. Drawing from cases of neurobiological research on potentiation phenomena, I argue that attempting to explain a phenomenon may provide reason to suspend judgment about its characterization, but this cannot provide warrant to recharacterize it if researchers cannot infer a phenomenon's characteristics from this explanation. To explicate this, I go beyond explanation – mechanistic or otherwise – to analyze why and how researchers change their epistemic commitments in light of new evidence.

Keywords Scientific Phenomena · Explanation · Representation · Epistemology of science

1 Introduction

A substantial body of literature in the philosophy of science identifies *scientific phenomena* as a target of scientific reasoning. A phenomenon is taken to consist in causal interactions. Given that they are stable and repeatable, scientists aim to characterize phenomena in order to explain why they acquire the empirical results they do in

✉ David Colaço
djc60@pitt.edu

¹ Department of History and Philosophy of Science, University of Pittsburgh, 1101 Cathedral of Learning, 4200 Fifth Avenue, Pittsburgh, PA 15260, USA

studies and to theorize about what is common to each of their occurrences (Bogen and Woodward 1988; Woodward 1989). Several popular accounts of scientific explanation now also take phenomena (so understood) to be what scientists aim to explain. This includes mechanistic explanation (Machamer et al. 2000; Glennan 2002; Craver 2007; Bechtel and Richardson 2010), whose adherents accept as a “platitude” that “each mechanism has a phenomenon” (Garson 2017, 104). These accounts propose, roughly, that a phenomenon is explained by schematizing its underlying causal structure. For these accounts, “to characterize a phenomenon correctly and completely is a crucial step” for formulating an “acceptable mechanistic explanation” (Craver 2007, 128).

Characterizing phenomena is important to achieve practical aims as well. Phenomena are causally efficacious; their occurrences have measurable effects (Hacking 1983). In fields like biology, where researchers aim to develop therapies for disease, researchers want to exploit the occurrences of phenomena to control the systems in which they occur (Craver 2007, 1). Determining the characteristics of a phenomenon allows researchers to determine how inducing it affects other components of the system. Researchers aim to determine what *causal role* a phenomenon plays in the circumstances in which it occurs, by determining what can cause this phenomenon to occur and what effects its occurrence can have.

Despite the importance of phenomena, until recently, little has been said about how their characterizations are *evaluated*. This is no minor concern: much is at stake for researchers when it comes to determining the adequacy of their target phenomenon’s characterization. Though it reflects researchers’ aims, researchers must evaluate to what extent this characterization is consistent with empirical evidence. If they do not revise this characterization in light of evidence that indicates that it is deficient, they risk dooming their research project (see Colaço 2018).

To their credit, the philosophers who endorse mechanistic explanation have addressed the recharacterization of phenomena. While these “mechanists” accept that characterizing a phenomenon is necessary for explaining it, many also argue that a phenomenon can be recharacterized based on its explanation. At first glance, scientific practice seems consistent with this argument. As the proliferation of memory phenomena suggests (Squire 2009), schematizing distinct mechanistic explanations for what was initially characterized as a single phenomenon can lead to its recharacterization. Thus, there is a question regarding why it is often the case that phenomena – the targets of explanation – are recharacterized following attempts to explain them. The mechanists suggest that this is due to the “feedback” between explaining and characterizing: mechanistic explanations provide insight into the adequacy of the target phenomenon’s characterization (Bechtel and Richardson 2010, 238).

I disagree. In order to determine what role explaining phenomena can play in recharacterizing them, I address a more basic question: *when should researchers recharacterize phenomena?* With the stakes when evaluating a phenomenon’s characterization in mind, I explicate why researchers may not recharacterize phenomena based on how they explain them. This is because, even when researchers aim to mechanistically explain a phenomenon, schematizing an explanation need not warrant recharacterizing it. That being said, explaining a phenomenon is not entirely irrelevant to evaluating its characterization. When researchers devise explanations of a phenomenon that raise concerns about the adequacy of its characterization, they *suspend judgment*, rather than reject it. Making this distinction helps to explain that not all

evidence relevant to critically evaluating a phenomenon's characterization warrant the same reaction from researchers: some findings *indicate* deficiencies and thus warrant revision or rejection, while others merely *suggest* deficiencies and thus warrant reexamination. Thus, my analysis here takes steps towards a thorough and more nuanced account of how scientific commitments change in light of new evidence.

To defend my position, I present a case study of the neural phenomenon long-term potentiation (LTP), a popular case amongst the mechanists. I introduce a mechanist account of characterizing phenomena in Section 2. In Section 3, I address the evaluation of a phenomenon's characterization, and I identify the difference between recharacterizing a phenomenon and suspending judgment about it. In Section 4, I discuss how LTP was initially characterized. I discuss how its characterization was evaluated following attempts to explain LTP in Section 5. In Section 6, I explain why some evidence provides warrant to recharacterize a phenomenon, while other evidence provides reason to suspend judgment about it. I show that evaluating a phenomenon's characterization requires researchers to directly determine the adequacy of the characteristics that they specify. I conclude by examining the limited sense in which explaining a phenomenon could provide warrant for recharacterizing it.

2 Characterizing a phenomenon

To characterize a phenomenon, researchers identify a set of co-occurring features and the conditions under which these features occur. They formulate a representation of these features along with these conditions, which I refer to as the *characterization of a phenomenon*. An account of this process comes from Craver and Darden.¹

According to Craver and Darden, researchers characterize a phenomenon to develop a “description of the behavior or product of the mechanism as a whole” (8), where mechanisms are “entities and activities organized such that they are productive of regular changes from start or set-up to finish or termination conditions” (Machamer et al. 2000). Importance is placed on characterizing phenomena in the project of mechanistic explanation, as “to discover a mechanism, one must first identify a phenomenon to explain” (Craver and Kaplan 2014, 275). A phenomenon's characterization “prunes the space of possible mechanisms”: it constrains the search for *mechanistic details*, which causally underwrite this phenomenon's occurrence (52). The mechanists distinguish between characterizing a phenomenon and explaining it. If one characterizes a phenomenon, one need not be able to explain how it occurs, or why it occurs under certain conditions. For mechanists, explaining a phenomenon consists in schematizing what causally underwrites its occurrence. Because they are explanatory and not descriptive, mechanistic details are not included in a phenomenon's characterization.

Craver and Darden take characterizing a phenomenon to have two parts. First, researchers determine the features that constitute the phenomenon of interest (the *features* of a phenomenon). Second, researchers establish the conditions under which these features occur (the *conditions* of a phenomenon). This includes precipitating

¹ Unless otherwise specified, all references are to Craver and Darden 2013.

conditions, which are “all of the many sets of conditions sufficient to make the phenomenon come about” (56), and inhibiting conditions, which are “the conditions under which the phenomenon fails or is blocked from occurring” (57). Characterizing phenomena in this way is not unique to the mechanists. Other accounts of explanation (Woodward 2003) and modeling (Batterman 2009) appeal to phenomena (so understood). All of these philosophers take phenomena to have “stable recurrent features which can be produced regularly by some manageably small set of factors” (Woodward 1989, 395). These qualities of phenomena make them “potential objects of explanation and prediction by general theory” (Woodward 1989, 393).

On Craver and Darden’s account, “a purported phenomenon might be recharacterized or discarded entirely as one learns more about the underlying mechanisms” (62). This is because the discovery of mechanistic details can reveal one of two errors: either there has been a “lumping together [of] two separable phenomena produced by different mechanisms” or a “splitting [of] one phenomenon into many,” where a phenomenon is mischaracterized as multiple phenomena (60). In either case, these details are *discrepant*: there is not a single characterization of a phenomenon and a single mechanistic explanation for its occurrence. On this account, a discrepancy between the characterized phenomenon and its mechanistic explanation should result in the phenomenon’s recharacterization. *Recharacterization* occurs when a phenomenon’s characterization is rejected, revised or changed to specify different characteristics, and subsequently accepted in this new form. This includes *splitting recharacterizations*, in which a characterization is rejected and subsequently fractured in its revision, resulting in characterizations of distinct phenomena.

The idea that an explanation can result in recharacterizing the explanandum phenomenon is found elsewhere in the mechanist literature. Craver suggests that “dissociating realizers mandates splitting kinds,” suggesting that a phenomenon’s characterization should be split if its manifestations are explained by distinct mechanisms (Craver 2004, 962). Craver and Darden’s account echoes Craver’s norms of explanation, which suggest that “lumping errors” are resolved when phenomena initially thought to be unitary are shown to be “composed of a variety of distinct and dissociable processes” (2007, 124). Other mechanists like Bechtel and Richardson claim that phenomena should be “reconstituted” following their explanation, as “our conception of what needs explaining... is shaped by the explanations and the models we develop” (2010, 194). They argue that mechanistic explanations require that the explananda phenomena “must themselves be understood” in different terms (Bechtel and Richardson 2010, 239). In addition, they argue that “there is feedback between phenomena and explanatory models” (Bechtel and Richardson 2010, 238). Thus, several mechanists have suggested that an explanation provides warrant for recharacterizing its explanandum, with some suggesting that discrepant mechanistic details provide warrant to split this explanandum. While mechanists acknowledge that practices beyond a mere feedback between explaining and characterizing might inform the recharacterization of a phenomenon – Craver, for example, discusses interlevel experiments where mechanistic components are manipulated to measure changes in the explanandum phenomenon (2007, 145) – they only address circumstances in which researchers assess and recharacterize phenomena for reasons related to the discovery of mechanistic details.

An alleged example of this is the split of procedural and declarative memory following research on patient H.M. After a surgery that removed his hippocampus, H.M. could not form new memories of events. However, H.M. could learn new tasks – he could form new procedural memories – though he would have no explicit memory of them (Scoville and Milner 1957). Mechanists suggest that H.M. provided evidence that “the mechanisms for procedural memory and declarative memory are distinct,” and treating memory as a unitary phenomenon amounts to “lumping together what we now know to be two distinct kinds of memory” (Craver 2004, 961). Along with other research, the study of H.M. “seemingly removed any doubt” that there are multiple memory phenomena, which ought to be characterized differently (Craver 2004, 962).

As a summation of the mechanistic perspective on scientific phenomena, Craver and Darden’s account has several virtues. They correctly assert that characterizing a phenomenon includes specifying the conditions under which it can be induced or inhibited, as well as specifying what the phenomenon itself is like. Determining a phenomenon’s features and how it is induced or inhibited are both needed to determine its causal role. Craver and Darden are also right that characterizing a phenomenon is distinct from explaining it, as explanatory questions about the phenomenon – why questions and how questions – cannot be answered with its characterization. And, most importantly for this article, they rightly identify that phenomena are often recharacterized following their mechanistic explanation.

Indeed, these mechanists might be onto something. If what we characterize as manifestations of the same phenomenon turn out to be underwritten by distinct causal structures, it seems reasonable to think that these structures may cue us into differences that should lead us to question whether we ought to characterize them as distinct phenomena to be theorized about and controlled differently. But, does this mean that mechanistic details provide warrant for recharacterizing explananda phenomena? To answer this question, we must go beyond phenomena as they relate to mechanisms and instead analyze their characterizations.

3 Epistemic attitudes towards a phenomenon’s characterization

To investigate when a phenomenon should be recharacterized, I first discuss the reasons why these characterizations are accepted. As mentioned in the introduction, philosophers have rightly identified that researchers characterize phenomena for theoretical and practical aims, even though they also may aim to mechanistically explain these phenomena. To fulfill theoretical and practical aims, researchers seek to accept accurate characterizations of phenomena. A phenomenon’s characterization is accurate to the extent that (1) the features specified in this characterization are consistent with the features that co-occur, and (2) the conditions specified to precipitate or inhibit the characterized phenomenon are consistent with the conditions under which the occurrence is observed. Accuracy in this sense is obtainable even though a phenomenon’s characterization is abstracted away from many of the details of its occurrences. This is because this phenomenon can potentially occur in different contexts, so long as these contexts include the conditions of its occurrence that are specified in its characterization. In other words, because a phenomenon is characterized to represent what is common to each of its occurrences, the idiosyncratic characteristics of each of its

manifestations are not represented. All that is needed is a characterization that describes what features will occur under a specified set of conditions. With an accurate characterization, researchers can explain and predict what happens when the phenomenon manifests, reason about the characteristics every manifestation has, and control its occurrences. An inaccurate characterization of a phenomenon will not explain or predict what consistently occurs, and it will also lead to ineffective interventions, limiting researchers' control over the target system.

Researchers assess the accuracy of a phenomenon's characterization in light of evidence related to the phenomenon's constitutive features, along with evidence related to the conditions under which it occurs. If the evidence for a phenomenon's characterization is consistent with what is specified in this characterization, researchers should accept it, and theorize and experiment in a way that is consistent with the characterized phenomenon's occurrence.² If more evidence is produced, then researchers should reassess their acceptance of this characterization.

It is not necessary that researchers either accept or reject a phenomenon's characterization upon reassessment, as there are more than two epistemic attitudes one could take towards it. There is a third kind of attitude, if 'attitude' is the right word for it: upon assessment of evidence, researchers may *suspend judgment* about a phenomenon's characterization.³ Suspending judgment involves more than merely neither accepting nor rejecting, as it seems wrong to say that researchers suspend judgment about claims that have not been considered (Friedman 2013, 168). Rather, judgment is suspended when researchers consider a characterization of a phenomenon and cannot determine what other attitude to adopt towards it. Instead, they adopt a "neutral state" towards this characterization (Friedman 2017, 307).

One might question the relevance of the difference between rejecting a characterization of a phenomenon and suspending judgment about it, given that, whatever other differences there are between them, they both amount to not accepting it. The answer lies in what researchers *do* in light of holding these respective attitudes. When deliberating on a phenomenon's characterization, a researcher "reflects on his evidence (or lack thereof) for and against" the characterization (Friedman 2013, 179). But, if the researcher "doesn't take his evidence to settle the matter," then they suspend judgment (Friedman 2013, 179). This is not where the inquiry ends, however. Suspending judgment has a "push towards its own demise: in suspending we ask a question and (at least in some minimal sense) seek an answer" (Friedman 2017, 316). This is because researchers need to characterize phenomena to achieve their aims. Thus, researchers want to resolve their suspended judgment. To do this, they perform more research.

Herein lies the difference between rejecting and suspending judgment. When researchers suspend judgment, they continue to inquire about a phenomenon's existing characterization. If this occurs, researchers investigate the putative phenomenon further, to evaluate its characterization in light of more definitive evidence. When they reject

² This is an idealization; researchers do not have strict criteria when they have sufficient evidence to accept a phenomenon's characterization, nor is there a sharp distinction between accurate and inaccurate characterizations. Nevertheless, researchers converge upon a judgment of what evidence is sufficient, and what degree of accuracy is desired.

³ These epistemic attitudes can be conceived of as continuous rather than as a trichotomy. Something like tweaking could be thought of as a practice based on an attitude that lies between suspending judgment and rejection.

this characterization, they move on to new ones, having reason to think that it is not accurate and therefore cannot be “saved” in its existing formulation. It is the latter that is the first step of recharacterization: rejecting a phenomenon’s existing characterization as deficient. Thus, while, in practice, researchers may not explicitly adopt epistemic attitudes about phenomena, we can learn about researchers’ reasoning from their response to acquiring new evidence.

With an understanding of differences between rejecting and suspending judgment, the nature of the disagreement between the mechanists and myself becomes salient. If researchers should reject a phenomenon’s characterization based on their explanation of it, then mechanistic details provide warrant for recharacterization. If they instead should suspend judgment, then, whatever epistemic role mechanistic details play, they do not provide this warrant. Mechanists suggest the former: researchers should reject a phenomenon’s characterization if they discover that distinct mechanisms underwrite its manifestations. This constitutes a “lumping error.” Further, they should revise its characterization, so that they have characterizations of distinct phenomena, each of which is consistent with a distinct mechanistic explanation. As these together constitute recharacterizing a phenomenon, this suggests that mechanistic details not only provide warrant to reject a phenomenon’s existing characterization but also to revise it. If this position is correct, then researchers should recharacterize a phenomenon in light of what they learn from explaining it, following which they should theorize and experiment in a way that is consistent with the recharacterized phenomenon’s occurrence.

My position, by contrast, is that the distinction between rejecting and suspending judgment provides a way of understanding how explanation can be relevant to evaluating the characterization of the phenomenon to be explained, without committing to the idea that explanation provides warrant to recharacterize it. When researchers have the theoretical and practical aims that I describe, I expect them to perform additional studies to evaluate the accuracy of a phenomenon’s characterization in light of discovering discrepant mechanistic details. This is because these discoveries typically do not provide warrant to reject these characterizations, let alone to recharacterize them.

I defend my position by analyzing a case of neurobiological research in which researchers inquired about a phenomenon’s characterization both prior to and following the discovery of discrepant mechanistic details. This case, long-term potentiation (LTP), illustrates how researchers react differently to evidence about a phenomenon’s characterization when compared to how they react to this phenomenon’s discrepant mechanistic details. Consistently throughout this case, we find researchers recharacterizing the phenomenon only when they produce evidence that is inconsistent with what is specified in their characterization of this phenomenon. By contrast, researchers reexamine but do not revise their phenomenon’s characterization in light of discrepant mechanistic details, which is best explained by the idea that researchers suspend judgment about this characterization in light of these discrepancies. This is because discrepant mechanistic details are merely suggestive of issues with the phenomenon’s characterization. This not only undermines the mechanists’ claims, but it also provides the basis for determining, more systematically, what provides warrant to recharacterize a phenomenon.

4 Formulating and evaluating a phenomenon's characterization

To show how my position captures how researchers evaluate characterizations of scientific phenomena, I describe the initial characterization of LTP. LTP is a neurobiological phenomenon that involves the potentiation of a neural synapse that lasts on a timescale of minutes or more. When a synapse is potentiated, there is an increase in the strength and number of excitatory responses of postsynaptic neurons following the activation of the presynaptic neurons. Its discovery is credited to Terje Lømo (Craver 2003).

Lømo demonstrated that the stimulation of presynaptic neurons in the hippocampi of anesthetized rabbits resulted in minutes-long potentiation of postsynaptic neuron populations (1966).⁴ He applied high-voltage inputs to the presynaptic cells in quick succession. Changes in the responses of the population of postsynaptic cells were recorded following stimulation. The intensity and number of excitatory postsynaptic responses increased when the presynaptic cells were stimulated, while the latency of the responses shortened. Lømo described the excitation of the presynaptic cell population, stimulated by his electrode, as a precipitating condition of this potentiation phenomenon. He also described how varying stimulation affects the intensity of the potentiation, as well as the time needed for the system to return to its initial state. Together, his specification of the experimental conditions that induced the potentiation, along with the description of the potentiation itself, consisted in the initial characterization of the phenomenon.

Following its formulation, this characterization of LTP was tested to determine its accuracy. In collaboration with Tim Bliss, Lømo determined the strength of the stimulation required to induce the effect (1973). Lømo and Bliss modified the experimental protocol they used to induce potentiation in order to understand the effect of the placement of the stimulating electrode. The researchers found no evidence that LTP resulted from these aspects of the experiment. This gave researchers a clearer sense of when LTP occurs. In distinct experiments with Tony Gardner-Medwin, Bliss performed experiments that were identical to Lømo's, except for the fact that the rabbits were not anesthetized. They observed the features of LTP "without anesthesia and under the more stable conditions," thereby precipitating the phenomenon in a way unspecified by Lømo initially (Bliss and Gardner-Medwin 1973, 358). Thus, Bliss and Gardner-Medwin provided evidence for the accuracy of Lømo's characterization of LTP – Lømo correctly identified some of conditions of its occurrence – but also identified other conditions that precipitate the phenomenon. With these findings, the researchers revised the characterization of LTP to convey that the phenomenon occurs under "approximately normal physiological conditions," specifying that LTP can be precipitated in physiologically-active organisms (Bliss and Gardner-Medwin 1973, 358). This revision occurred without schematizing a mechanistic explanation for LTP. At the time, Lømo, Bliss, and Gardner-Medwin did not attempt to explain LTP. This indicates that there are dimensions to evaluating a phenomenon's characterization that are independent from attempting to explain it, which is absent from mechanist accounts of recharacterizing a phenomenon.

⁴ Research on LTP was motivated by research – such as studies on H.M. – that suggested that the hippocampus is responsible for memory formation. However, Lømo himself did not test the relation between LTP and memory (Lømo 2003).

The work from Lømo, Bliss, and Gardner-Medwin informed what I call the *standard characterization* of LTP, which includes precipitating and inhibiting conditions. Thus, initial research on characterizing LTP took the following trajectory. Lømo formulated a characterization of a phenomenon. Bliss and Gardner-Medwin induced the phenomenon under different but consistent conditions. Researchers tested what conditions were needed to precipitate the features specified in LTP's characterization: their findings provided reason to accept that Lømo's initial characterization was accurate in the sense that none of features and conditions that Lømo specified failed to occur when the phenomenon was induced, though it could be made more detailed. Overall, in this initial phase, evaluating LTP's characterization amounted to the direct test of the specified characteristics of this phenomenon.

5 Recharacterization and explanation

The previous section shows how researchers use evidence about a phenomenon's features and the conditions under which it occurs to formulate its characterization. But, what about explanation? In this section, I address the evaluation of LTP's characterization following attempts to mechanistically explain it. In these attempts, two kinds of discrepancies were discovered. One discrepancy relates to the phases of LTP putatively explained by different activities; the other relates to manifestations called 'LTP' throughout the nervous system putatively explained by the activity of different receptors. I contrast the evidential role of empirical findings related to the features of LTP and the conditions under which it occurs with findings related to mechanistic details in these two cases.

The first discrepancy stems from the discovery of the molecular details of LTP in the region investigated by Lømo starting in the late 1970s (see Teyler and DiScenna 1987), after which the role of the NMDA glutamate receptor in LTP was experimentally demonstrated (Collingridge et al. 1983).⁵ Following these discoveries, hippocampal LTP was thought to be explained by the increase of glutamate receptors (Collingridge 1985). As a result of stimulation, more glutamate is released into the synaptic cleft, which results in AMPA and NMDA glutamate receptors reacting (Collingridge and Bliss 1987, 289). The discrepancy relates to the activities of these receptors. Distinct mechanistic schemata can be devised to explain the so-called "phases" of LTP, inciting an ongoing debate regarding "whether LTP is divided into temporal phases, each involving different biochemical interactions" (Huang 1998, R350). The activation of NMDA receptors sets off a chain of reactions that bring additional AMPA receptors to the dendrites of the postsynaptic cells. This results in the potentiation of the neuron, called early-phase LTP. If stimulation continues to occur, the postsynaptic cell will undergo changes in transcription enzymes, which result in the production of AMPA receptors, called late-phase LTP. Thus, researchers discovered that distinct sets of mechanistic details underwrite what were thought to be all manifestations of LTP, and the sets do not always occur in tandem.

⁵ LTP research continued in earnest, as researchers faced difficulties developing experimental paradigms to investigate the molecular basis of this phenomenon (Nicoll 2017, 281).

A second discrepancy was discovered concurrent with research on the phases of LTP. In areas other than the hippocampus, and thus outside of the purview of Lømo's investigations, potentiation effects that were initially thought to be consistent with the standard characterization of LTP were discovered. However, it was found that these potentiation effects are underwritten by an inconsistent set of mechanistic details, some of which do not involve the NMDA receptor (see Sullivan 2017 for a historical review). By the end of the 1980s, the question of whether or not studies of what was thought to be LTP in the nervous system are "related to hippocampal LTP" became the focus of reviews of the state of LTP research (Teyler and DiScenna 1987, 150; see also Collingridge and Bliss 1987, 288). Calling the standard characterization "NMDAR-dependent LTP," researchers aimed to determine whether or not potentiation effects in other parts of the nervous system are adequately described by this characterization. Again, different mechanistic details were found to underlie what were initially thought to be manifestations of the same phenomenon in different areas of the nervous system.

Given that these discoveries suggest that different manifestations of LTP are best explained according to distinct mechanistic schemata, we can examine what role these discoveries played in evaluating LTP's characterization. If the discovery of discrepant mechanistic details provides warrant for recharacterizing a phenomenon, we should expect that both of these discrepancies provided warrant to recharacterize LTP. In other words, we should find that researchers characterize the phases and the occurrences of LTP in other parts of the nervous system as distinct phenomena based on the discovery of these discrepancies.

However, this is not what has occurred. Despite agreement that discrepant mechanistic details were discovered – there was "actually little disagreement on the truly key experimental results" because "the experiments were easily reproducible" (Nicoll 2017, 284) – as I will show, there was disagreement about whether (and if so how) LTP should be recharacterized. While there have been different responses to these discrepancies, we find the same reasoning at play.

Related to the first discrepancy, the differences between the mechanistic details associated with the respective phases were investigated in order to determine if there are differences between the features of LTP or the conditions of its manifestations when LTP is underwritten by one set of mechanistic details as opposed to the other. Several researchers have argued that it is problematic to "ignore [the] possible distinction" between LTP's phases (Huang 1998, R350). For example, Emily Huang argues that "the induction requirements of the two phases... differ," and the potentiation persists to different lengths of time (1998, R350). Huang suggests treating "LTP as a dynamic phenomenon that may be modulated by many factors, including time," highlighting the importance of identifying LTP's precipitating conditions when evaluating its characterization (1998, R352).

Alternate conclusions have been drawn regarding LTP's phases. Bliss and colleagues suggest that characterizing LTP in terms of phases that have different timespans is misguided: they argue that early-phase LTP is of "variable duration," so it cannot be differentiated from late-phase LTP in this way (2018, A109). Instead, they argue that "the critical factor that determines whether the potentiation comprises" one or both phases of LTP "is the timing (and potentially also the strength) of the induction trigger" (Bliss et al. 2018, A109). Thus, researchers again identify LTP's precipitating conditions when evaluating its characterization. Bliss and colleagues also suggest that "the

existence of two potentially long-lasting forms of LTP can explain numerous conflicting data on the transduction and expression mechanisms of LTP” (2018, A109). They argue that LTP’s phases – what they call LTP1 and LTP2 – have characteristics that might elucidate their mechanistic schemata. This is not what one expects of a lumping error. Researchers do not suggest recharacterizing LTP in terms of its mechanistic details; rather, they suggest reschematizing these details based on LTP’s characteristics.

These positions contrast the fact that, before data had been collected on the differences between the phases of hippocampal LTP, researchers did not suggest recharacterizing the phenomenon. Bliss and Collingridge previously argued that the relation between hippocampal potentiation phenomena “has not been clearly defined,” so any distinction made between LTP’s phases must be done “tentatively” (1993, 34). Huang, Bliss, and Collingridge’s claims illustrate that debate over the phases of LTP is one about the differences between the characteristics of instances counted as LTP.

Related to the second discrepancy, research on potentiation that occurred in different neural regions led researchers to argue that LTP’s standard characterization does not accurately characterize the features or conditions that were measured in all of the contexts in which potentiation was induced. J. Sweatt remarks about the “different types of LTP in the mammalian CNS – hippocampal, cortical, cerebellar, NMDA receptor-dependent and independent – just to name a few prominent categories” (1999, 399). Other researchers concur, arguing that there are forms of LTP that “may share some, but certainly not all, of the properties and mechanisms of NMDAR-dependent LTP” (Malenka and Bear 2004, 5). In their review, Robert Malenka and Mark Bear state that it “is now clear that LTP” is not a “unitary” phenomenon; they recommend that researchers “define at which specific synapses these phenomena are being studied... and how they are being triggered,” highlighting the differences between the characteristics of what were initially counted as manifestations of the same phenomenon (2004, 5). Jay Blundon and Stanislav Zakharenko echo this sentiment, arguing that “since the first characterization by Lomo [sic] and Bliss, use-dependent potentiations have been described in other parts of the central nervous system,” but it “quickly became clear that these LTPs may differ (2008, 599).⁶

Thus, following the discovery of the second kind of discrepancy, we find some indication of a convergence amongst researchers that sufficient evidence has been put forward to recharacterize LTP. Sweatt, Malenka and Bear, and Blundon and Zakharenko all argue that LTP should be fractured into characterizations of distinct phenomena that occur throughout the brain. These distinct phenomena are explained in terms of distinct mechanistic schemata: Malenka and Bear, for example, argue that some forms of LTP do not involve the NMDA receptor (2004). However, we do not find the discovery of discrepant mechanistic details alone providing reason to recharacterize LTP. We see that researchers routinely address the induction of LTP, claiming that various phenomena once called ‘LTP’ have different precipitating conditions. More recently, we have also seen researchers point out that the stimulation that triggers NMDAR-dependent LTP serves as an inhibiting condition for other potentiation phenomena (Kullmann and Lamsa 2008). Researchers also identify that the

⁶ They also note, “the brain region, the neuron type in that region, and types of inputs that synapse on a particular neuron type are all major determinants of the type of LTP” (Blundon and Zakharenko 2008, 599).

potentiation strength varies between these phenomena, illustrating the differences between their features (see also Abraham 2003).

While this section does not exhaust the history of LTP research following Lømo, it illustrates the relation between explaining and characterizing this phenomenon. The discovery of both discrepancies led researchers to evaluate how they had characterized LTP and investigate the manifestations of what were initially characterized as the same phenomenon. For the first discrepancy, there remains debate about whether the phases of LTP ought to be characterized as distinct phenomena or if the phenomenon should be characterized as dynamic in nature. For the second discrepancy, many researchers now argue that there are distinct potentiation phenomena that should be characterized as such. These recharacterizations were, in many cases, no mere tweaks: the features and conditions were both mischaracterized in different cases. Some revisions were radical, reformulating precipitating conditions into inhibiting ones and vice versa.

In both cases, however, the reasoning was the same. Researchers discovered discrepant mechanistic details, which led them to reexamine its once-accepted characterization. Contemporary discussions of the characterization of LTP reflect this reasoning: debates regarding whether the phases of LTP are the same phenomenon, are parts of a single phenomenon, or are distinct phenomena are in part *motivated* by the differences between the mechanistic details that underwrite each phase, but these debates are not *resolved* by the discovery of these details. Instead, researchers aim to resolve these debates by investigating (1) the conditions that precipitate or inhibit the potentiation, and (2) the persistence and strength of the potentiation that occurs. These are the phenomenon's conditions and features, respectively.

6 Recharacterizing a phenomenon

The case of LTP – starting from the formulation of its characterization, to attempts to explain it in terms of the mechanisms that underwrite its occurrence – illustrates how a scientific phenomenon's characterization is evaluated. To draw normative implications from this example, I reflect on how the evaluation of a phenomenon's characterization relates to the theoretical and practical aims of researchers who characterize it. I examine why researchers typically reexamine, rather than reject, their characterizations in light of the discovery of discrepant mechanistic details. By aligning a phenomenon's characterization with the researchers' aims, I show that my conclusions are not unique to the case of LTP.

6.1 Reasons to recharacterize a phenomenon

The discovery of mechanistic details preceded some researchers pushing for the splitting recharacterization of LTP, but this fact fails to capture what provided warrant for this recharacterization. Following each discovery, researchers performed additional experiments; with the findings from these experiments, they evaluated their characterization of LTP. With this case in mind, a clearer picture of when phenomena should be recharacterized can be discerned.

Based on the actions of the researchers, what provided warrant to recharacterize LTP? In its evaluation following attempts at its explanation, LTP was recharacterized

based on evidence produced in experiments that had one of the following forms: researchers either observed the features of LTP under conditions that are inconsistent with those specified in the characterization of LTP, or they produced the conditions of LTP and observed features that are different from those specified in the characterization of LTP. Throughout Section 5, we saw interest in the induction of LTP's phases and the conditions that precipitate LTP in different areas of the brain, and we saw interest in systematic differences between the strength of potentiation phenomena. The results of these kinds of experiments relate directly to the components of a phenomenon's characterization. They serve as a test of the specified characteristics of the phenomenon.

More generally, when researchers have theoretical and practical aims, a phenomenon should be recharacterized when the descriptions researchers make about its occurrences based on what is specified in its characterization are not consistent with the evidence that they collect: either inconsistent features occur under the specified conditions of the phenomenon's occurrence, or the features occur under conditions that are inconsistent from those that are specified. It is under these circumstances that a phenomenon's characterization should be rejected. If this evidence provides reason to think that these specified characteristics should be changed, so that the new characterization is consistent with the available evidence, then a phenomenon's characterization should be revised and accepted in its new form. Combined, this is when a phenomenon should be recharacterized.

What role did the discovery of mechanistic details play in recharacterizing LTP? The discoveries that resulted from attempts to explain LTP led researchers to reexamine the standard characterization. In Section 5, we saw that the activities of distinct receptors involved in instances called LTP led researchers to search for differences between the characteristics of these instances. What accounts for this reexamination of a phenomenon's characterization? The answer lies in the fact that differences between the entities or activities of distinct mechanistic details suggests that their respective initiations might be sensitive to different external causes or that different external effects might be sensitive to their respective terminations. If either of these possibilities were the case, then manifestations of the phenomenon in question underwritten by distinct mechanistic details would have different characteristics.

Why does the discovery of discrepant mechanistic details provide reason to reexamine the characterization of a phenomenon they explain but not provide warrant for its recharacterization? The answer is based on the limitations of what can be inferred about phenomena from the mechanistic details that underwrite them. The fact that researchers determine that distinct mechanistic details underwrite instances characterized as the same phenomenon means that there may be contexts in which the characterization inaccurately describes the occurrence under investigation. However, the discovery of discrepant mechanistic details is not in and of itself inconsistent with a single characterization of the phenomenon. This is because distinct mechanisms may be initiated by the same conditions or may produce the same outcome. Given that, for theoretical and practical aims, a characterization of a phenomenon *only* describes the features of a phenomenon and the conditions under which it occurs, the mechanistic details that underwrite it may make no difference to its characterization.

This is why it is useful to distinguish between suspending judgment and rejecting a phenomenon's characterization. Some findings provide reason to think that a

phenomenon's characterization is deficient from the perspective of their aims – in the LTP case, deficient in terms of accuracy – and therefore should be rejected or even recharacterized. Other findings are merely suggestive: they provide reason to suspend judgment and evaluate a phenomenon's characterization, to provide warrant to either accept or reject it. This distinction is illustrated in how Bliss and colleagues react to LTP1 and LTP2: their tentativeness in 1993 contrasts their confidence in 2018 following further research on LTP's characteristics. Thus, explaining a phenomenon can lead to recharacterizing it, but pointing this out does not identify the evidential nature of this relation. Even if, as a matter of historical fact, findings that are suggestive of issues with a phenomenon's characterization often precede research that reveals that it is deficient, this need not be the case. So long as it is possible to acquire more evidence, researchers can suspend judgment about scientific claims, and inquire into their adequacy. Following this inquiry, researchers can acquire evidence that provides warrant to recharacterize a phenomenon.

6.2 The aims of characterizing a phenomenon

My position that researchers suspend judgment about a phenomenon's characterization following the discovery of discrepant mechanistic details captures what researchers investigating LTP actually did. This history is best explained by the fact that suspending judgment about a phenomenon's characterization in this kind of situation is more in line with researchers' aims when compared to simply rejecting it. In other words, if researchers have the theoretical and practical aims that I have described, then the discovery of mechanistic details alone typically will not provide warrant to recharacterize their target phenomenon.

Suspending judgment but not rejecting a phenomenon's characterization in light of discrepant mechanistic details aligns well with researchers' theoretical aims.⁷ If researchers aim to develop a general theory about a target system when investigating "objects of research" (Feest 2017) – for instance, Bliss and colleagues argue that splitting LTP1 and LTP2 based on their characteristics can inform the theorization of the neural basis of memory (2018, A109) – researchers identify the phenomena present in the system under investigation and determine whether this theory predicts that the phenomenon as characterized will occur. If discrepant mechanistic details are discovered, researchers must investigate whether the differences between the mechanistic details are relevant to the causal role of what they each underwrite. However, from a theoretical point of view, making a typological distinction between phenomena with the same conditions and features – and thus play the same causal role – is counterproductive, even if instances are mechanistically explained differently. This is because the aim is to theorize about a type-level characterization of what constitutes the phenomenon as well as what it can do in the target system. The ability to theorize about this type-level characterization is undermined if researchers split this characterization along mechanistic lines, even though the phenomenon in question plays the same causal role.

⁷ The aim of developing characterizations and explanations of phenomena into a theory that "can be used together to describe, predict, explain, and test aspects" of a target of investigation is also an idea shared by mechanists (Craver 2007, 175).

The fact that researchers suspend judgment rather than reject a phenomenon's characterization in light of the discovery of discrepant mechanistic details also aligns with their practical aims. Phenomena are characterized so that they can be exploited to induce changes in the systems in which they occur. To induce the desired changes, researchers must determine what conditions precipitate or inhibit the phenomenon of interest and what effects the phenomenon can have. Characterizing as different what can play the same causal role is counterproductive to inducing changes to fulfill these practical aims. Thus, from a practical perspective, a phenomenon should only be recharacterized to reflect differences in features of the phenomenon and conditions under which it occurs, as these differences are what make a difference when exploiting the phenomenon's occurrence to cause other changes in the system.

The point of all of this is not that these are the only aims that researchers could have, nor is it that there could be no other criteria for characterizing phenomena. Rather, what is important is that these aims are common and intuitive for researchers to hold, and they are fulfilled via accuracy.⁸ In addition, these aims are held concurrently with researchers' aims to mechanistically explain the phenomenon of interest, which shows that theoretical and practical aims need not be in direct competition with mechanistic explanatory aims. Given that explanation typically does not provide evidence about an explanandum phenomenon's characteristics, it is not because of evidence of mechanistic details that researchers recharacterize phenomena.

6.3 What would it take for explanation to provide warrant for recharacterization?

I have argued that mechanistic details typically do not provide warrant to recharacterize a phenomenon, at least for the commonly-shared aims that I have described. What would need to be the case for mechanistic details to provide this warrant? The answer is that mechanistic details would provide warrant to recharacterize an explanandum phenomenon if the *characteristics of the schematized mechanisms* could be independently demonstrated to be inconsistent with its specified features or the conditions of this phenomenon's occurrence.

This claim rests on characterizing a mechanism's components as well as its set-up and termination conditions. In establishing these set-up and termination conditions, researchers may determine that their characterization of the conditions that precipitate the occurrence of the phenomenon are inconsistent with the conditions that initiate the mechanism that causally underwrites the phenomenon and may determine how to better characterize these conditions in light of what is known about the mechanism's initiation conditions. For instance, researchers might discover that the conditions that initiate one of the mechanisms cannot concurrently occur with the precipitating conditions specified in the characterization of the phenomenon. This would provide evidence inconsistent with a phenomenon's characterization, which would provide reason to reject it, given the differences between what initiates each mechanism and the precipitating conditions that are specified in its existing characterization.

⁸ LTP is not the only example of recharacterization in the way that I describe. The case of H.M. is not one in which a phenomenon's characterization was split based on mechanistic explanations. This is because, short of determining that procedural memory did not involve the activity of the hippocampus, the distinction of the two was determined before either were explained. It seems more plausible that procedural and declarative memory were split because the phenomena associated with them were distinguished from one another.

However, pointing out this way in which mechanistic details could provide warrant for recharacterizing phenomena reveals why it is not typically the case that they do so. The discovery of discrepant mechanistic details need not tell researchers about the relationship between the mechanism's input and the conditions specified in the phenomenon's characterization, nor need it tell researchers about the relationship between the mechanism's output and the phenomenon's features. Furthermore, rejecting an existing characterization of a phenomenon resolves only part of the inquiry. To *recharacterize* a phenomenon, researchers must also revise this existing characterization. The fact that distinct mechanistic details causally underwrite a single characterized phenomenon does not provide evidence for how the characterization should be revised, unless it also provides details about how researchers should modify their specification of its features and conditions.

One might object to this conclusion by arguing that differences between mechanistic details entail that they must explain distinct phenomena. If one were committed to the view that every phenomenon is underwritten by a unique mechanism, then schematizing distinct mechanisms would be sufficient for recharacterizing the phenomenon they underwrite, or at least would be sufficient for rejecting its existing characterization.

This objection can be read in two ways, and neither reading is defensible. First, this objection could suggest that phenomena should be differentiated whenever distinct mechanisms causally underwrite them. However, it is unclear why such a commitment is warranted. Researchers can determine the causal role phenomena play, and how they can exploit their causal properties, fulfilling their theoretical and practical aims. These aims are not fulfilled if researchers treat instances with the same causal role as distinct phenomena. Second, the objection could suggest that, if there are differences between mechanisms, then the instances that are characterized as the same phenomenon must be different in ways relevant to their characterization. However, that there is a difference does not entail that this difference must be relevant to counting something as a manifestation of a phenomenon. This is because a phenomenon's characterization describes what set of conditions and features are shared amongst its manifestations. The shared features of the phenomenon may not vary if distinct mechanisms underlie them. This is why researchers identify the shared characteristics of a phenomenon's manifestations: what underwrites the specified features that co-occur under the specified conditions is not relevant unless they make a difference to their co-occurrence or causal role.

What if one objects in the other direction? Craver has expressed conventionalist leanings on the topic, suggesting that "human perspective enters into decisions about which mechanisms matter for splitting or lumping a putative kind" (2009, 585). He also argues that "there is no objectively appropriate degree of abstraction for typing mechanisms" (Craver 2009, 589). If one were committed to this view, then convention is what matters both when determining whether mechanistic details provide warrant for recharacterizing explananda phenomena and when determining whether mechanistic details are discrepant with this characterization.

My emphasis on theoretical and practical aims should make it clear that I do not deny that convention matters when characterizing phenomena. These aims are "conventions" – popular conventions, which favor accuracy – but conventions nonetheless. That being said, it is important to recognize that the conventions regarding how phenomena are characterized need not be determined by the conventions regarding

how they are explained. As we saw in Section 5, Bliss and colleagues seem to share the explanatory aims of LTP researchers, but they still suggest reschematizing mechanistic schemata based on novel evidence about LTP's characteristics. This is because characterizing a phenomenon is not always, let alone solely, done in the effort of explaining it. This is why there are cases, like LTP, where there is a characterization of a unitary phenomenon according to certain conventions, but it is explained (according to different conventions) by distinct mechanistic schemata that constitutively underlie its respective manifestations. In these cases, there is nothing about the conventions of the researchers that should lead us to expect that these explanations will provide warrant to recharacterize the explanandum.

This fact supports the fundamental thesis of this paper. Characterizing phenomena is an important scientific process, which has its own conventions. Its various aims cannot be properly understood, let alone satisfied, solely from the perspective of explanation. While characterizing and explaining phenomena are linked in the ways that I have discussed in this paper, a satisfactory account of why phenomena are recharacterized cannot be obtained without investigating this process unto itself. This is what is missing in the philosophical accounts according to which phenomena are the targets of theory and control, which is why the literature on the topic can benefit from a reorientation from solely discussing how phenomena are explained and modeled to discussing how they are characterized in the first place.

7 Conclusion

I have provided an account of how characterizations of phenomena are evaluated in light of evidence related to the empirical commitments specified in these characterizations. In doing so, I have identified the sense in which phenomena can be recharacterized following the discovery of mechanistic details that causally underwrite the phenomenon's manifestations. What is in question is what role these details play. I have argued that the discovery of discrepant mechanistic details typically does not provide warrant for recharacterizing its explanandum phenomenon. Occurrences with different schematized mechanistic explanations can be characterized as manifestations of the same phenomenon if the differences between these schematized mechanisms do not result in differences between the features of the phenomenon, or the conditions under which the features occur. Whether there are differences and what those differences are only can be determined with additional studies.

Nevertheless, the discovery of mechanistic details can provide reason to suspend judgment and reexamine a phenomenon's characterization: the differences between the mechanisms may relate to differences between the features of a phenomenon, or the conditions under which it occurs. This conclusion provides more general insight into when researchers should suspend judgment about empirical claims. When researchers have evidence that suggests that their claims are not adequate, they should suspend judgment and inquire into them. This inquiry is achieved by doing more research.

I do not deny that schematizing the mechanistic details that causally underwrite a phenomenon's occurrence is a viable way to explain it. Rather, I point out that determining a discrepancy in its mechanistic details typically does not tell researchers the right things about a phenomenon to warrant recharacterization. Evidence about the

features of a phenomenon and the conditions under which it occurs is what matters to evaluating its characterization. This makes clear how researchers evaluate a phenomenon's characterization, which many philosophers, mechanistic and otherwise, view as integral to theorization, explanation, and control.

Acknowledgements Thanks to William Bechtel, Liam Kofi Bright, Mazviita Chirimuuta, Eric Hochstein, Edouard Machery, Kenneth Schaffner, Jacqueline Sullivan, and James Woodward for comments on previous versions of this paper. A draft of this paper was presented at the 2018 meeting of the Eastern Division of the American Philosophical Association. Thanks to those who provided me feedback at this session, and to Sara Aronowitz for her commentary. This paper was also presented as part of the Neural Mechanisms Webconference. Thanks to those who provided me feedback during this web session.

References

- Abraham, W. C. (2003). How long will long-term potentiation last? *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 358(1432), 735–744. The Royal Society.
- Batterman, R. W. (2009). Idealization and modeling. *Synthese*, 169(3), 427–446.
- Bechtel, W., & Richardson, R. C. (2010). *Discovering complexity: Decomposition and localization as strategies in scientific research*. Cambridge: MIT.
- Bliss, T. V. P., & Collingridge, G. L. (1993). A synaptic model of memory: Long-term potentiation in the hippocampus. *Nature*, 361(6407), 31.
- Bliss, T. V. P., & Gardner-Medwin, A. R. (1973). Long-lasting potentiation of synaptic transmission in the dentate area of the unanaesthetized rabbit following stimulation of the perforant path. *The Journal of Physiology*, 232(2), 357–374.
- Bliss, T. V. P., & Lømo, T. (1973). Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the Perforant path. *The Journal of Physiology*, 232(2), 331–356.
- Bliss, T. V. P., Collingridge, G. L., Morris, R. G. M., & Reymann, K. G. (2018). Long-term potentiation in the hippocampus: Discovery, mechanisms and function. *Neuroforum*, 24(3), A103–A120.
- Blundon, J. A., & Zakharenko, S. S. (2008). Dissecting the components of long-term potentiation. *The Neuroscientist*, 14(6), 598–608.
- Bogen, J., & Woodward, J. (1988). Saving the phenomena. *The Philosophical Review*, 97(3), 303–352.
- Colaço, D. (2018). Rip it up and start again: The rejection of a characterization of a phenomenon. *Studies in History and Philosophy of Science Part A*, 72, 32–40.
- Collingridge, G. L. (1985). Long term potentiation in the hippocampus: Mechanisms of initiation and modulation by neurotransmitters. *Trends in Pharmacological Sciences*, 6(1985), 407–411.
- Collingridge, G. L., & Bliss, T. V. P. (1987). NMDA receptors-their role in long-term potentiation. *Trends in Neurosciences*, 10(7), 288–293.
- Collingridge, G. L., Kehl, S. J., & McLennan, H. (1983). Excitatory amino acids in synaptic transmission in the Schaffer collateral-commissural pathway of the rat Hippocampus. *The Journal of Physiology*, 334(1), 33–46.
- Craver, C. F. (2003). The making of a memory mechanism. *Journal of the History of Biology*, 36, 153–195.
- Craver, C. F. (2004). Dissociable realization and kind splitting. *Philosophy of Science*, 71(5), 960–971.
- Craver, C. F. (2007). *Explaining the brain: Mechanisms and the mosaic Unity of neuroscience*. New York: Oxford University.
- Craver, C. F. (2009). Mechanisms and natural kinds. *Philosophical Psychology*, 22(5), 575–594.
- Craver, C. F., & Darden, L. (2013). *In search of mechanisms: Discoveries across the life sciences*.
- Craver, C. F., & Kaplan, D. (2014). Towards a mechanistic philosophy of neuroscience. In S. French & J. Saatsi (Eds.), *The Bloomsbury companion to the philosophy of science* (pp. 268–292). London: Bloomsbury Academic.
- Feest, U. (2017). Phenomena and objects of research in the cognitive and behavioral sciences. *Philosophy of Science*, 84(5), 1165–1176.
- Friedman, J. (2013). Suspended judgment. *Philosophical Studies*, 162(2), 165–181.
- Friedman, J. (2017). Why suspend judging? *Noûs*, 51(2), 302–326.

- Garson, J. (2017). Mechanisms, phenomena, and functions. In *The Routledge handbook of mechanisms and mechanical philosophy* (pp. 122–133). Routledge.
- Glennan, S. (2002). Rethinking mechanistic explanation. *Philosophy of Science*, 69(S3), S342–S353.
- Hacking, I. (1983). *Representing and intervening: Introductory topics in the philosophy of natural science*. Cambridge: Cambridge University Press.
- Huang, E. P. (1998). Synaptic plasticity: Going through phases with LTP. *Current Biology*, 8(10), R350–R352.
- Kullmann, D. M., & Lamsa, K. (2008). Roles of distinct glutamate receptors in induction of anti-Hebbian long-term potentiation. *The Journal of Physiology*, 586(6), 1481–1486.
- Lømo, T. (1966). Frequency potentiation of excitatory synaptic activity in the dentate area of the hippocampal formation. *Acta Physiologica Scandinavica*, 68(Suppl 277), 128.
- Lømo, T. (2003). The discovery of long-term potentiation. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 358(1432), 617–620.
- Machamer, P., Darden, L., & Craver, C. F. (2000). Thinking about mechanisms. *Philosophy of Science*, 67(1), 1–25.
- Malenka, R. C., & Bear, M. F. (2004). LTP and LTD: An embarrassment of riches. *Neuron*, 44(1), 5–21.
- Nicoll, R. A. (2017). A brief history of long-term potentiation. *Neuron*, 93(2), 281–290.
- Scoville, W. B., & Milner, B. (1957). Loss of recent memory after bilateral hippocampal lesions. *Journal of neurology, neurosurgery, and psychiatry*, 20(1), 11.
- Squire, L. R. (2009). Memory and brain systems: 1969–2009. *Journal of Neuroscience*, 29(41), 12711–12716.
- Sullivan, J. A. (2017). Long-term potentiation: One kind or many? In M. Adams, Z. Biener, U. Feest, & J. Sullivan (Eds.), *Eppur Si Muove: Doing history and philosophy of science with Peter Machamer; a collection of essays in honor of Peter Machamer* (pp. 127–140). New York: Springer International Publishing.
- Sweatt, J. D. (1999). Toward a molecular explanation for long-term potentiation. *Learning & Memory*, 6(5), 399–416.
- Teyler, T. J., & DiScenna, P. (1987). Long-term potentiation. *Annual review of neuroscience*, 10(1), 131–161.
- Woodward, J. (1989). Data and phenomena. *Synthese*, 79(3), 393–472.
- Woodward, J. (2003). *Making things happen: A theory of causal explanation*. New York: Oxford University.

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