

Constitutive Relevance in Interlevel Experiments

Maria Şerban*

Sune Holm[†]

Abstract

One reason for the popularity of Craver’s mutual manipulability account of constitutive relevance (MM) is that it seems to make good sense of the experimental practices and constitutive reasoning in the life sciences. Two recent papers (Baumgartner and Gebharder [2016]; Baumgartner and Casini [2017]) propose a theoretical alternative to (MM) in light of several important conceptual objections. Their alternative approach, the No De-Coupling (NDC) account conceives of constitution as a dependence relation which, once postulated, provides the best explanation of the impossibility of breaking the common cause coupling of a macro-level mechanism and its micro-level components. This entails an abductive view of constitutive inference. Proponents of the NDC or abductive account recognize that their discussion leaves open a big question concerning the practical dimension of the notion of constitutive relevance: Is it possible to faithfully reconstruct constitutional reasoning in science in terms of a failure to de-couple, via interlevel experiments, phenomena from their mechanistic constituents? Focusing on the field of memory and LTP research, this paper argues that the abductive account provides a more adequate description of interlevel experiments in neuroscience. We also suggest that the account highlights some significant practical recommendations of how to interpret the findings of interlevel experiments.

Keywords: constitution; mechanism; interventionism; abductive inference; inter-level experiments

*Philosophy Section, Department of Media, Cognition, and Communication, University of Copenhagen, Karen Blixens Vej 4, 2300 Copenhagen, Denmark
mariaserban.org@gmail.com

[†]Philosophy Section, Department of Media, Cognition, and Communication, University of Copenhagen, Karen Blixens Vej 4, 2300 Copenhagen, Denmark
suneh@hum.ku.dk

1	Introduction	2
2	Mutual Manipulability and Constitutive Relevance	4
2.1	Constitutive relevance through an interventionist lens	4
2.2	Mutual manipulability: A methodological application	7
3	Trouble for the Mutual Manipulability Account	10
4	The Abductive Account of Constitution	12
5	The Abductive Account: A Methodological Application	14
5.1	LTP and memory experiments	15
5.2	A comparative summary	22
6	Conclusions	24

1 Introduction

The success of the mechanistic outlook in philosophy of biology is in no small part due to its capacity to describe and analyze the discovery and reasoning strategies commonly used in the life sciences. This paper takes issue with the adequacy of one prominent mechanistic approach that targets constitutive reasoning and the associated experimental practices of biologists. Craver’s mutual manipulability account (MM) of constitutive relevance (Craver [2002], [2007b]) has been previously criticized for introducing an internal tension between an interventionist analysis of causation and the standard non-causal view of mechanistic constitution (Leuridan [2012]; Fagan [2013]; Schindler [2013]). One such recent criticism has led to an alternative approach ‘which allows for treating constitution and causation as two closely related, but still distinct, types of dependencies.’ (Baumgartner and Gebharder [2016], p. 753; see also Baumgartner and Casini [2017]) However, proponents of this alternative view of mechanistic constitution acknowledge that further philosophical investigation is required to show whether the ensuing account can guide a faithful reconstruction of constitutive experimentation and reasoning in the life sciences. We take on this challenge and investigate the implications that the abductive or No-Decoupling (NDC) theory of mechanistic constitution (Baumgartner and Casini [2017])¹ has for understanding the structure of interlevel mechanistic experiments and the inferences grounded in them.

¹The two names used by Baumgartner and Casini [2017] to designate their alternative to Craver [2007b] mutual manipulability account capture different, yet related, aspects of their analysis of constitution and constitutive inference. We shall explicate these two features in more detail in sections 4 and 5, but we follow these authors in using both names in the paper.

Methodological and explanatory considerations always seem to go hand in hand in a mechanistic framework. To wit, it has been argued that experimental and modelling practices in the life sciences aim to identify and explore mechanisms underlying particular biological phenomena, and therefore that biological explanations describe the mechanisms that produce these target phenomena. More specifically, mechanistic explanations are taken to describe how an organised collection of active mechanistic parts interact with each other to produce some systemic (macro-level) property or behaviour in a complex causal system. This standard characterization of mechanistic explanations captures their dual character: they are causal descriptions that span multiple levels of mechanistic organization. Some have argued that this points to the possibility of distinguishing two types of mechanistic explanations being used in biology (Ylikoski [2013]; Fagan [2013], [2015]). Causal mechanistic explanations describe how some system-level effect is produced by the organised workings of the micro-level mechanistic parts, whereas constitutive mechanistic explanations specify how the macro-level complex causal system is built from its micro-level active constituent parts. But even those who maintain that the mechanistic framework covers two distinct types of explanation seem to agree that the discovery or experimental practices underlying these explanations are the same (Fagan [2015]). That is, causal relevance experiments and constitutive relevance experiments have a similar structure and are appropriately analysed in the same conceptual framework. This assumption is also at the core of Craver’s mutual manipulability (MM) approach to mechanistic constitution (Craver [2007b]).

Craver’s approach is based on the idea that the interventionist framework, intended to analyse the structure of causal reasoning and inference in science (Woodward [2003]), can be extended to characterise interlevel mechanistic experiments and the constitutive inferences based on such experiments. He claims that ‘[t]he norms of constitutive relevance are implicit in the experimental strategies that neuroscientists [and other scientists] use to test claims about componency and in the rules by which [they] evaluate instances of those strategies’ (Craver [2007b], p. 144). In other words, the mutual manipulability account of constitutive relevance is supposed to offer a descriptively adequate picture of the interlevel mechanistic experiments carried out in the life sciences. We challenge the widespread agreement that Craver is right about the structure of interlevel (or constitutive relevance) mechanistic experiments and the type of inferences they are able to ground.

Our strategy is both critical and constructive. To start with, we pick up on the objections articulated in two recent papers by Baumgartner and collaborators (Baumgartner and Gebharder [2016]; Baumgartner and Casini [2017]) summarising the conceptual changes required to build a consistent and empirically serviceable approach to analyze constitutive reasoning in biology. The positive spin of this criticism is the abductive view of constitution that conceives of constitution as a *sui-generis* dependence relation which

is required to explain the impossibility of experimentally breaking the common cause coupling of a macro-level mechanism and its micro-level components.

Our main focus in this paper is on the structure of interlevel mechanistic experiments and the inferences based on them. The plan is this. In section 2 we review the key features of the mutual manipulability approach that are meant to establish the distinction between the relations of constitutive dependence (or relevance) and causal dependence. Then, in section 3 we use Baumgartner’s and Gebharder’s [2016] critical arguments to uncover the vulnerabilities of the mutual manipulability account. We introduce the conceptual scaffolding of Baumgartner’s and Casini’s [2017] abductive theory of mechanistic constitution in section 4, which we then expand in section 5 to sketch an alternative analysis of interlevel mechanistic experiments and their epistemic roles. As the main case study we will use an example that has become a *locus classicus* in the mechanistic literature, namely the case of LTP and memory experimentation. In doing this, we take on the challenge of showing how the abductive account can be used to provide an adequate reconstruction of mechanistic constitutive reasoning in the biological sciences (Baumgartner and Gebharder [2016], p. 753; Baumgartner and Casini [2017], p. 19). More generally, this methodologically-applied analysis will help to sharpen the distinction between causal and constitutive experiment-based inferences by clarifying the justificatory structure of interlevel mechanistic experiments.

2 Mutual Manipulability and Constitutive Relevance

This section offers a brief summary of Craver’s mutual manipulability (MM) account of mechanistic constitutive relevance. Although the approach has been the object of scrutiny for many previous studies, we think that revisiting some of these common places will facilitate the critical task to follow.² Our reconstruction also seeks to set on equal ground the two pillars of Craver’s approach to the problem of constitutive relevance: the conceptual scaffolding of interventionism and its commitment to descriptive adequacy with respect to the experimental methodologies of the life sciences.

2.1 Constitutive relevance through an interventionist lens

The primary motivation behind Craver’s account of constitutive relevance is to elaborate a consistent notion of mechanistic explanation that would also mirror the experimental

²Authors adopting the MM account primarily for its methodological benefits include: Illari and Williamson [2012]; Kaplan [2012], and Fagan [2013]. On the other hand, Harbecke [2010], [2015]; Leuridan [2012]; Schindler [2013] among others have discussed challenges facing the MM account of constitutive relevance within a mechanistic framework.

methodologies of practicing biologists (Craver [2002], [2007a], [2007b]). Although Craver is an explicit endorser of the ontic view of scientific explanation, it is easier to understand the mutual manipulability approach to constitutive relevance from a more methodological or epistemic perspective. So we will adopt this stance throughout the paper.³

According to a widespread consensus, mechanisms are complex causal systems of multiple components whose activities and interactions are organized across multiple levels to produce some overall phenomenon (cf. Illari and Williamson [2012]; Fagan [2015]). Thus, within a mechanistic framework the problem of constitutive relevance is tightly related to the problem of individuating or demarcating what counts as a mechanistic level. A mechanistic level is individuated by the entities and activities that are organised so as to produce a detectable or recognisable phenomenon. This generic definition suggests that experimental techniques play an important role in distinguishing one mechanistic level from another. Subsequently they should also figure (or be taken into account) in defining the criteria for constitutive relevance in a mechanistic framework. (We will return to this point in section 2.2.)⁴

Craver characterises the constitution relation linking higher mechanistic levels to lower ones in terms of four distinctive features.⁵ Firstly, constitution is a part-whole relation in the sense that mechanistic constituents at the lower level overlap spatiotemporally with the broader mechanism at the higher level they are part of. The second feature restricts this sense of spatio-temporal congruence. It requires that proper mechanistic components (or constituents) make a difference, under some appropriate conditions, to the behaviour of the mechanism as a whole. Thirdly, constitutive dependence is a relation that goes both ways (from the whole to its mechanistic parts and *vice versa*). Therefore, some type of supervenience relation is supposed to hold between any two mechanistic levels. That is, any change in what counts as the higher mechanistic level must be associated with some change in at least one of the lower level mechanistic constituents.⁶ Finally, most mechanists follow Craver and Bechtel [2007] in claiming that constitution is a distinct dependence relation than causation. In support of this distinctiveness claim they invoke the atemporality and the bidirectionality of the constitutive dependence relation.⁷

This introduction of the core features of constitutive relevance paints a picture of ‘stacked’ or ‘nested’ mechanistic levels. There is talk of higher and lower levels which can also be substituted conveniently with talk of macro and micro-levels (typically when the

³This will also facilitate our discussion of the arguments presented in Baumgartner and Gebharter [2016] and Baumgartner and Casini [2017] who take a similar stance.

⁴A further consequence of focusing on the experimental methodologies used to identify and delimit mechanistic levels is that one can only settle empirically questions such as: how many mechanistic levels are there? And, what are the (relevant) constituents of such levels?

⁵See Craver [2007b] for a detailed discussion of these features.

⁶Baumgartner and Gebharter [2016], pp. 735-7 offer a more detailed discussion of the supervenience criterion. See also Harbecke [2014].

⁷For more details, see Craver [2002]; Craver and Bechtel [2007]; Bechtel [2017]; and for a particularly poignant critique of the distinction, see Leuridan [2012].

relationship between only two mechanistic levels is considered). We follow Craver [2007b] and illustrate these notions by appeal to neuroscientific research on the mechanism of long-term potentiation (LTP).

LTP is a heavily studied mechanistic model of how information is stored at the cellular-molecular level in the brain and of how neural change may contribute to spatial learning and memory. In a mechanistic framework, one can think about the relationship between spatial memory and the mechanism of LTP in terms of at least three different levels: the behavioural-organismic level, the electrical-synaptic level and the molecular-kinetic level.⁸ The behavioural level is associated with different types of learning and memory and the conditions under which they are achieved, the electrical synaptic level focuses on the special properties of the generation and propagation of action potentials, while the molecular level corresponds to the biochemical properties and reactions of various neuroreceptors, ions and enzymes. Constitutive relevance relations can be searched for and established between any two of these mechanistic levels and as we shall see in more detail in section 5, different types of experiments target these more specific dependence relations.

How is experimental LTP research to be understood within the MM framework? At the core of Craver's approach is the claim that some working spatio-temporal part of a mechanism is constitutively relevant to the mechanism as a whole if and only if one can manipulate both the part to bring about a change in the whole and manipulate the whole so as to produce a change in the part. For example, some acting microlevel entity (for example, the NMDA receptors generating Ca^{2+} channels through the membrane) is constitutively relevant to the macrolevel phenomenon (the macroscopically measured LTP response) if and only if the former is a spatiotemporal part of the latter and it is possible to perform one bottom-up and one top-down intervention that will produce changes that are detectable at the two-levels. MM is meant to provide both necessary and sufficient conditions for constitution relations: if the acting macrolevel mechanism and its working spatio-temporal part are mutually manipulable then the latter is constitutively relevant with respect to the former; if the two levels are not mutually manipulable, then the constitutive relevance relation does not hold. The notion of manipulation used here is supposed to reflect closely the experimental practices encountered in the life sciences and to be operationalisable in the conceptual framework of interventionism (Woodward [2003]).

MM is best viewed as an extension of interventionism from a theory of causal dependence relations (and causal inference) to a theory that covers both causal and non-causal constitutive relations. Woodward [2003] analyses causal dependence relations in terms of

⁸More mechanistic levels can be distinguished empirically, but our discussion of the experimental methodologies used in LTP research will focus primarily on these 3 levels. For another classification see Craver [2002], [2007b].

the notion of intervention which broadly speaking is an idealized experiment that eliminates all confounding factors to isolate the relation between two variables: X and Y . That is, an intervention I on a variable X with respect to another variable Y is a causal process that guarantees that any change in the value of Y is due to a change in the value of X . It follows then that X is a cause of Y if and only if the relation between X and Y is invariant under some intervention on the value of X in the sense that changing the value of X determines a change in the value of Y . As such, interventions model causal relevance experiments. The MM account contends that the notion of an intervention can be used to model interlevel or constitutive relations as well. This is supposed to work by extending the notion of intervention to cover processes that show both the dependence of macrolevel variable on the microlevel one and *vice versa*.

In summary, according to MM, a constitutive relation is established if and only if one can determine spatio-temporal contiguity between the macro and microlevel variables and the possibility of two types of interventions, one bottom-up and one top-down which show the mutual dependence of the value of the macrolevel variable on the microlevel one and *vice versa*. Before discussing the conceptual troubles encountered by the MM approach, we would like to make the methodological pillar of MM even more apparent.

2.2 Mutual manipulability: A methodological application

A prima facie strength of Craver's approach derives from its empirical adequacy with respect to scientific practices. That is, the MM approach is taken to provide an accurate and philosophically insightful picture into the experimental methodology by which mechanistic models are constructed in the life sciences. Since we will later take issue with this claim, it is worth discussing more explicitly how Craver links this experimental methodology to the demarcation of mechanistic levels and thus to the MM account of constitutive relevance.

To this purpose, we focus next on the notion of a mechanistic experiment. In a nutshell, to perform a generic mechanistic experiment is 'to place one or more parts of the mechanism under the experimenter's control and to assess from its responses to such interventions how the experimental results can be used to prune and shape the construction, evaluation, and revision of mechanism schemas' (Craver and Darden [2013], p. 299; see also Craver [2002]). Mechanistic experiments can test either for causal relevance, or for constitutive relevance or they can be complex combinations of experiments which address 'particular questions' about candidate entities and activities within a mechanistic schema. All types of experiments can be analysed in terms of three basic elements: the experimental model, the intervention technique and the detection technique being used. We will focus on the difference between causal and componenty or interlevel experiments.

Causal relevance experiments 'provide evidence about what makes a difference to what

is in a mechanism. Such experiments are used to determine the start or setup conditions for a mechanism and to establish the active organization among its components, that is, to establish how something happening at one stage of the mechanism produces, inhibits, modulates, or otherwise makes a difference to what happens at other stages.’ (Craver and Darden [2013], p. 315) Interlevel experiments, on the other hand, are designed to test constitutive relevance and so they ‘are used to answer two related questions. First, one might ask whether a given part is relevant to a given explanandum phenomenon at a higher level, that is, to the behaviour of a mechanism as a whole. Second, starting with a higher-level phenomenon, one might wonder which parts in the system are relevant to the phenomenon (that is, which parts are components in the mechanism for that phenomenon).’ (ibid.)⁹

Within the class of interlevel experiments, Craver and Darden [2013] further identify two types of bottom-up experiments, interference experiments and stimulation experiments, and one type of top-down experiments, namely activation experiments. As their name suggests, bottom-up experiments presuppose that the researcher intervenes on a component in a mechanism and detects changes in the behaviour of the mechanism as a whole, whereas in top-down experiments, the intervention is ‘localized’ at the macro-level (the whole mechanism level) whereas detection targets changes in the activities or properties of the components of the mechanisms. Craver and Darden [2013] contend that determining constitutive relevance relations requires both bottom-up and top-down experiments (and even mixed-type experiments) to be performed in the appropriate research context.¹⁰

The link between the previous taxonomy of mechanistic experiments and the MM approach should be obvious by now. MM says that constitutive relevance relations are conclusively established by two types of interventions which allegedly map onto the bottom-up and top-down types of inter-level experiments used in biological research. Returning to the LTP example, the two parts of the MM criterion correspond to the conjunction of top-down experiments that, for instance, record electrical signals in the hippocampus during learning tasks, and of bottom-up experiments such as synapse saturation and behavioural pharmacological experiments (see Section 5).

The alignment between the experimental methodology employed in this particular neurobiological research context and the MM account makes the latter an appealing

⁹Other quotes suggest that the difference between the methodology of the two types of experiments is not as big as this classification seems to suggest. For instance, Craver and Darden write that ‘interlevel experiments [...] are not different in kind from experiments for testing causal relevance, but rather should be viewed as particular kinds of such experiments’. (Craver and Darden [2013], p. 324)

¹⁰The epistemic work of these different kinds of experiments is sometimes characterised in terms of the epistemic notion of robustness, which is typically understood in terms of manipulation or detection of a phenomenon via multiple theoretically independent routes (Wimsat [1994]). For instance, Craver [2002] writes that: ‘Interlevel experiments are interesting because they provide different independent paths of access to a phenomenon that are nonetheless part of the same multilevel theory’. (Craver [2002], p. 95)

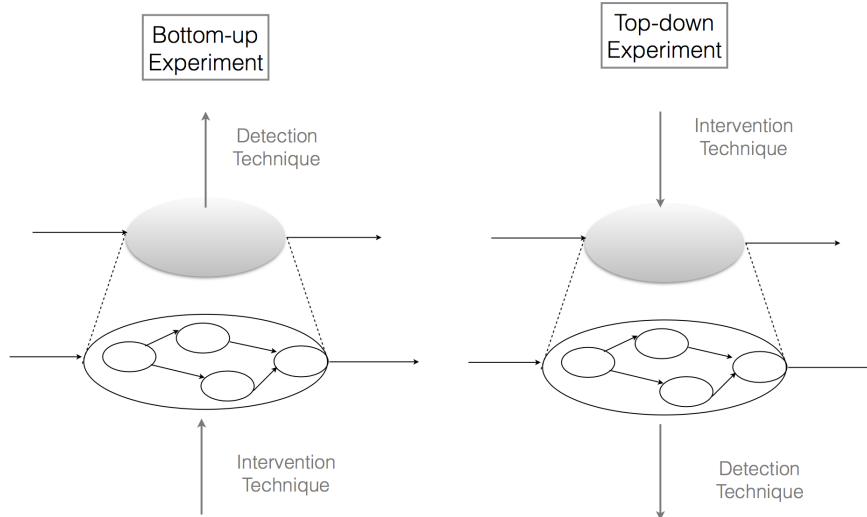


Figure 1: Mechanistic representation of interlevel experiments. Adapted from Craver and Darden [2013], p. 126.

philosophical approach to the problem of constitutive relevance. But is this really enough? Craver seems to assume that the bottom-up and top-down interlevel experiments are both necessary and sufficient for establishing constitutive relations between the macro and micro-level of a target mechanistic phenomenon. In other words, his approach entails that there is a direct inference from these types of experiments to a conclusion about LTP being a mechanism for information storage (and memory) at the cellular-molecular level. Just like conceiving of causal relevance experiments in terms of (ideal) interventions sheds light on causal inference and causal reasoning in science, just so thinking of interlevel experiments in terms of the mutual manipulability approach is meant to clarify the structure of constitutive inferences in mechanistic science.

We turn next to the evaluation of this claim. MM has been repeatedly challenged on grounds of not drawing a sharp enough boundary between constitution and causation (Leuridan [2012]), of ruling out the asymmetry of constitutive explanation (Schindler [2013]), and of falling short both as a normative constraint on constitutive explanations and as a descriptively adequate analysis of interlevel experiments (Harbecke [2010]). Two recent papers (Baumgartner and Gebharter [2016] and Baumgartner and Casini [2017]) have proposed a critical approach dealing with the weaknesses of the MM approach. Their authors have shown that repairing the conceptual scaffolding of MM entails a different analysis of the structure of constitutive inference and the associated interlevel experiments. In what follows, we summarise the core of their critical arguments and show that their alternative analysis of constitutive inference provides a more insightful account of the experimental methodologies used in LTP and memory research.

3 Trouble for the Mutual Manipulability Account

Above we have suggested that embedding the account of constitutive relevance into an interventionist framework spells trouble for MM. In this section we unpack a forceful argument articulated in Baumgartner and Gebharder [2016] and partially rehearsed in Baumgartner and Casini [2017] which analyses the problematic relationship between interventionism and the MM approach to constitutive relevance. We then highlight the amendments to MM that this criticism entails.

In mounting their argument, Baumgartner and Gebharder [2016] endorse two basic assumptions of Craver’s approach: (i) the non-causal character of the dependence relation of constitutive relevance and (ii) the formal definition of a ‘surgical’ intervention embedded in the double mutual manipulability criterion. The first leg of their argument assumes that it is possible to make an intervention on the macrolevel variable (let this be identified by the characteristic macrolevel behaviour ψ) with respect to each of its microlevel constituents (with characteristic behaviours ϕ_i). Given that an intervention is a cause acting at the level of the macrovariable (ψ), the change produced in the microlevel variable (ϕ_i) can occur on one of two causal pathways: (i) $I \rightarrow \psi \rightarrow \phi_i$ or (ii) $I \rightarrow \psi$ and $I \rightarrow \phi$. The first causal pathway is incompatible with the assumption that constitution is a different type of dependence relation than causation. The second causal pathway is however at odds with the formal definition of an intervention because it shows that I is a common cause for the macrolevel and microlevel variables. Ruling out both options amounts to showing that there is no viable constitution dependence relation as defined by the MM criterion.

On the second leg of their argument, Baumgartner and Gebharder [2016] propose a modified version of interventionism. Following Woodward [2015], they extend the notion of intervention to accommodate supervenience relations. Now I is an intervention* variable for X with respect to Y if and only if I causes X , I acts like a switch for all the other variables that cause X , and any causal pathway from I to Y goes either through X or through a variable Z on which X supervenes. In addition, I must be statistically independent of every cause of Y that is neither located on a path through X nor on a path through variable Z . A further modification guarantees that the effect of the intervention variable I on X occurs at time t before its effect on variable Y which can be detected at time $t' > t$.

By embedding MM in the modified account of interventionism*, the common cause model captured by the causal path (ii) above becomes viable. According to this model, the intervention variable is a common cause of the macrolevel mechanism and microlevel constituent entity. This gives rise to a different problem for the MM account of constitutive inference. Recall that according to MM, for the microlevel acting entity ϕ_i to be a constituent of the macrolevel mechanism ψ , there must be an intervention I that by




Figure 2: Representation of common-causal and constitutive models of mechanistic phenomena. Adapted from Baumgartner and Gebharder [2016], p. 750.

changing the value of ψ would also change the value of ϕ_i and there must also be an intervention on ϕ_i with respect to ψ that would change the value of the latter in virtue of changing the value of the former. Consider the first type of intervention which arguably models top-down experiments. The variable I_ψ is a cause of the macrolevel mechanism ψ and a cause of the microlevel working entity ϕ_i . Even if this does not entail that there must be an interlevel causal dependence relation (option blocked by modifying the definition of an intervention), MM is not able to guarantee a direct inference about ϕ_i being a constituent of the macrolevel mechanism ψ . The data produced by the top-down manipulation cannot be informative about the relationship between the variables ψ and ϕ_i . In other words, the non-surgical, fat-handed intervention can generate only confounded data because I_ψ is a common cause for the two target variables.

With the modifications entailed by interventionism*, Craver’s MM criterion cannot be used to distinguish empirically between a common-cause model and a constitutive model of mechanistic phenomena. That is, interlevel experiments do not rule in favour of the constitutive model (a) and against the common cause model (b). This entails a systematic underdetermination of the inference about the putative constitutive dependence relation between ψ and ϕ_i . Assuming that the micro and macrolevels are coupled via common causes which no interlevel experiment will be able to break leads to a different account of constitutive inference than the one implied by Craver’s account. This insight is further elaborated in Baumgartner’s and Casini’s [2017] abductive account of constitutive relevance to which we turn next.

4 The Abductive Account of Constitution

The criticism of MM sketched above implies that interlevel experiments are not quite what they seem at first. That is, contra Craver, Baumgartner and collaborators (Baumgartner and Gebharder [2016]; Baumgartner and Casini [2017]) claim that no amount of bottom-up or top-down mechanistic experiments would be able to establish directly and conclusively that a microlevel working entity is constitutively relevant for a macrolevel mechanism (phenomenon). Instead they propose that such experiments provide at best indirect evidence for the existence of constitutive relevance relations.¹¹

According to Baumgartner and Casini [2017], interlevel experiments reveal correlations holding between the values of the macrolevel and microlevel variables. They do so by manipulating simultaneously different mechanistic levels, on different causal paths. Conceived as fat-handed, non-surgical interventions*, top-down and bottom-up manipulations cannot break the inextricable coupling between phenomena and their mechanistic constituents. That is they cannot de-couple macrolevel mechanisms from their microlevel constituents.

As pointed out above, the evidence made available by fat-handed interventions is compatible with both a constitutive and a common cause model of the target mechanistic phenomenon. Comparing common cause models and constitutive models can reveal two facts. First, that interlevel experiments cannot be used to distinguish empirically the two models of mechanistic phenomena. And, second that the constitutive model is explanatorily superior to the common cause model because: ‘[it] not only reproduces the empirical correlations but also explains why the common-cause coupling of ψ and ϕ is not broken’. (Baumgartner and Casini [2017], p. 14)

The systematic failure of fat-handed interventions to de-couple macrolevel mechanisms from their microlevel constituents is what provides the indirect evidence for constitutive relevance relations. The resilient coupling of different mechanistic levels can be explained only by postulating that there is a *sui-generis* (non-causal) dependence relation connecting them. On the other hand, Baumgartner and Casini [2017] claim that experimental manipulations can be used to discover any surgical causes at macrolevel ψ which do not count as causes at the microlevel ϕ . Such manipulations would determine conclusively that there is no common-cause coupling between the two mechanistic levels. However, they insist that no amount of experimentation would be able to establish that a partic-

¹¹The notion of indirect evidence used by Baumgartner and Casini [2017] is relatively different from that discussed in other philosophical debates. Given theory T which entails hypotheses H_1 and H_2 , evidence e for H_2 indirectly supports H_1 for which there is (yet) no evidence available (Laudan and Leplin [1991]; Okasha [2002]; Wernndl [2013]). In the case at hand, H_1 and H_2 would correspond to the constitutive and common cause model, respectively. As we emphasize above, according to Baumgartner and Casini [2017] neuroscientific experimentation cannot provide direct evidence for H_1 , but systematic failures of interlevel experiments count as indirect evidence for H_1 which, unlike H_2 , has the benefit of being able to explain the impossibility of the breaking the common cause coupling between the mechanistic macro- and microlevel, respectively.

ular microlevel mechanistic element ϕ_i is not constitutively relevant for the macrolevel phenomenon. In other words, constitutive irrelevance of particular mechanistic microconstituents can also be established only abductively.¹²

This twofold consequence is at odds with Craver's original MM approach, for Baumgartner and Casini maintain that interlevel experiments can be used only to falsify conclusively the common-cause coupling between different mechanistic levels. Constitutive dependence (or independence) for them is entirely a matter of abductive inference. After performing multiple bottom-up and top-down experiments which show that two specific levels are inextricably connected (or coupled) via common causes can researchers infer abductively that the working entities at the lower mechanistic level are constitutively relevant for the higher level or for the phenomenon of interest.

If Craver's MM approach makes causal (relevance) inferences very similar to constitutive (relevance) inferences, Baumgartner's and Casini's alternative seems to do a better job at distinguishing the two types of inferences. They propose that '[t]he inference to constitution is inherently abductive: constitutional models are preferable over pure causal models because they explain both the highly correlated behaviour of phenomena and their constituents as well as the *impossibility to de-couple* them.' (Baumgartner and Casini [2017], p.15)

There are some striking differences in the pictures that the mutual manipulability and the abductive account paint of the purported constitutive experiments performed in neuroscience. Craver's MM account entails that the existence of an appropriate pair of top-down and bottom-up experiments can license constitutive inferences (to the effect that some microlevel working entities are constitutively relevant for some macrolevel mechanism). Baumgartner and Casini [2017] on the other hand contend that constitutive relevance relations can be conclusively established only after all possible interventions (both top-down and bottom-up) have been performed. Or, since such an experimental exploration of the space of all possible interventions is pragmatically impossible, constitutive inferences are at best supported abductively by experiments which fail to de-couple the mechanistic macrolevel from the microlevel.¹³

Establishing constitutive relations requires an extended test series exploring the whole space of possible ways of breaking the coupling of macro and micro levels. Only if these tests are unsuccessful, an inference to constitution is warranted. And since the evidence for the unbreakability of common cause couplings is never conclusive, constitution can only be inductively corroborated. (Baumgartner and Casini [2017], p. 18).

¹²'Overall, as to NDC, every constitutional inference – be it to constitution or to non-constitution – inevitably involves an inductive leap.' (Baumgartner and Casini [2017], p. 18)

¹³This further suggests a stronger asymmetry between constitution and causation concerning the type of underdetermination associated with establishing/verifying each type of relation.

Granting the internal consistency of the abductive (or non-decoupling) account, ‘[a]n obvious question remains: is it possible to faithfully reconstruct constitutional reasoning in science in terms of NDC?’¹⁴ This methodological test will not only vindicate NDC (or abductive) account as an alternative analysis of constitutive relevance but it will also help advance a number of questions in the philosophy of experimentation such as: Do causal and interlevel mechanistic experiments support different kinds of inferences? What types of evidence do mechanistic experiments provide? What is the relationship between explanation and confirmation in mechanistic research?

5 The Abductive Account: A Methodological Application

In this section we discuss the methodology of interlevel experiments in the framework of the abductive account. As a test case study¹⁵, we first introduce some paradigmatic experiments from LTP and memory research and then we apply the key ideas of the abductive account to clarify the structure and outputs of these experiments.

The analysis of interlevel experiments from the contemporary neurobiology of learning and memory is crucial for Craver’s account of the mosaic unity of neuroscience. That is, the MM account is supposed to yield insight on how multiple explanatory projects associated with different neurobiological experiments are integrated into more complete mechanistic explanations of target macrolevel phenomena. This view of the integrative unity of neuroscience has been challenged. Sullivan [2009] has questioned its descriptive adequacy with respect to the current experimental practices of neuroscience, while Levy [2016] has challenged its strength as a normative ideal for the field as a whole.

The abductive account seems to avoid these objections. First, the account does not imply that all experiments from LTP and memory research converge on their aim of building constitutive models or supporting constitutive inferences. In fact, an abductive view of constitutive reasoning in neuroscience is consistent with the idea that the majority of experiments carried out in this area of research are causal (or at least have a causal component). Second, instead of directly supporting a global integrative or reductive

¹⁴See also Baumgartner and Gebharder [2016], p.22.

¹⁵The philosophical analysis of the experiments carried out in the neurobiology of learning and memory can be seen as a ‘hard’ case study in the sense of Scholl and Raz [2016]. The existence of several conceptual frameworks (Craver [2007b]; Sullivan [2010]; Craver and Darden [2013]; Harbecke [2010]; Harbecke [2015]; Harbecke [2015b]; Gebharder [2017]) for analysing the structure and implications of these experiments testifies to their relevance for understanding the broader experimental landscape of neurobiology. Our aim is to show that the abductive account provides an insightful analysis of this case study which has some novel implications that distinguish it from the other accounts. However, we don’t claim that the abductive account will be the final word on the structure of experimentation in neuroscience. The categories introduced by the abductive account to analyse this case study must be further put to the test in discussing other experimental practices.

view of interlevel experiments, the abductive account is compatible with there being a multiplicity of experimental protocols which are tied to distinct explanatory goals pursued in the cognitive neurobiology of learning and memory (cf. Sullivan [2009]).

More specifically, the account predicts that interlevel experiments aim and consistently fail to break the common-cause coupling of macrolevel phenomena and their microlevel parts. Such consistent failures show not only that the correlations between the macrolevel and the microlevel phenomena are stable, but also that they are necessary. That is, they indirectly show that microlevel property or behaviour is necessary for the target macrolevel property or behaviour. While the abductive view remains silent with respect to the metaphysical distinction between constitutive and causal dependence relations, it manages to draw a sharper methodological distinction between the central characteristics and implications of purported constitutive and causal mechanistic experiments. In particular, it implies that the underdetermination problems facing constitutive experiments are different from those encountered in the case of causal experiments.

5.1 LTP and memory experiments

LTP is operationally defined as the abrupt and sustained increase in the efficiency of synaptic transmission following high-frequency (tetanic) stimulation of afferent fibers. The LTP response is almost always measured macroscopically across populations of synaptic contact. Macroscopic recordings are thought to track the *microscopic* LTP mechanisms which are the putative cellular information storage or memory mechanisms. Microscopic LTP involves the activation of NMDA receptors and of other types of glutamatergic receptors, intracellular Ca^{2+} and lipid cascades, retrograde messages, and other components, whose activities and interactions result in metabolic and morphologic alterations in the pre-synaptic and post-synaptic neurons. Some of the characteristics of microscopic LTP, including its rapid induction, persistence, and correlation with natural brain rhythms, have been taken to suggest a connection between microscopic LTP mechanisms and memory storage or learning.¹⁶

Different experimental methodologies are used to determine how LTP is related to the behavioural phenomena of spatial (and non-spatial) learning and memory. LTP was first detected *in vivo* in the hippocampus, a neural structure known to be activated during spatial learning and memory.¹⁷ Initial recordings of the electrical properties of this type

¹⁶As numerous reviews of the LTP literature reveal, there is quite some variation of theoretical perspectives concerning the higher level phenomenon that microscopic LTP is supposed to explain (Shors and Matzel [1997]; Holscher [2001]; Granger and Nicoll [2014]). The most prominent candidates are memory, cellular information storage, and learning. In the rest of the section we will focus on experimental research which links microscopic and macroscopic LTP to memory and/or learning. See also Sullivan [2010].

¹⁷For complementary historical perspectives on how LTP research came to be linked to learning and memory, see Holscher [2001]; Craver [2003]; Andersen *et al.* [2006].




Figure 3: Schematic representation of the main elements of LTP induction in a hippocampal neuron. Strong repetitive stimulation of the input to a hippocampal neuron can activate AMPA and NMDA receptors. The opening of the NMDA channel allows Ca^{2+} entry, which activates a particular protein kinase (CaM). The activation persists after the synaptic events that trigger it, and appears to be both necessary and sufficient for LTP induction in some neurons. The next series of steps that lead to the expression of LTP is still a matter of debate (Sanes and Lichtman [1999]; Granger and Nicoll [2014]).

of synaptic transmission were made by placing electrodes into two neural regions of the rabbit's brain: the perforant path and the dentate gyrus. Although these experiments allowed researchers to keep the brain intact as much as possible, there is a limit to the level of detail these detection devices can have access to. This led to the development of a series of more refined *in vitro* experimental methodologies which could yield information about the cellular and molecular elements involved in microscopic LTP. Although the hippocampal slice technique opens up experimental possibilities not available for *in vivo* preparations, the *in vitro* preparation requires cutting many of the existing neural connections.¹⁸

The mechanistic experiments carried out in LTP research yield information that is relevant for both causal and constitutive model-based inferences about the structure of the LTP response and its link to macrolevel behavioural phenomena of learning and memory. Offering a tidy classification of the variety of LTP experiments carried out in different laboratories and research centers goes beyond the more modest aims of this paper. While such a classification would further clarify the scope and limits of the applicability of the abductive account of constitutive inference, we hope that the selective sample of experiments discussed below would provide a clear enough illustration of the power of the account under scrutiny.

¹⁸This is just one among the many methodological challenges that make inferences from the *in vitro* experimental results to the relation between LTP and the hippocampus' function in memory so difficult to draw. For a more detailed discussion of the factors affecting the reliability of *in vitro* LTP experiments, see Shors and Matzel [1997]; Holscher [2001].

More specifically, we focus on those LTP experiments that seem to target constitutive relevance relations in the following methodological sense.¹⁹ They are interlevel experiments that involve either (i) experimental manipulations of the mechanism of LTP which are expected to bring about changes in behavioural phenomena such as memory formation or learning, or (ii) experimental manipulations of macrolevel behavioural phenomena which should determine changes in the LTP response. The first type of experiments correspond in the mechanistic taxonomy to bottom-up inter-level experiments whereas the second type qualifies as top-down inter-level experiments. Although for the purpose of testing the abductive account, analyzing the latter type of experiments would be more informative, there is a surprising scarcity of experiments that use behavioural manipulations of memory or learning tasks to study the properties of LTP (cf. Shors and Matzel [1997]; Holscher [2001]). We also aim to show that bottom-up experiments, far from being conclusive in establishing a necessary link between LTP and memory, are often supplemented by analogical inferences (based on the similarity of properties such as rapid induction, stimulus specificity, associativity and endurance, observed in the lower and higher level phenomena, cf. Gallistel [1995]; Dudai [1995]; Craver [2003]; Gallistel and Matzel [2013]; Trettenbrein [2016]). We take the latter inferential strategy to be compatible, if not directly predicted by the abductive account of constitutive inference.

Most LTP experiments involve the measurement of the excitatory postsynaptic response (EPSP) produced by the electrical stimulation of a neuron. ESPSPs can track single neuron or whole neural populations responses.²⁰ A reliable way for inducing LTP *in vivo* as well as *in vitro* is high-frequency stimulation (HFS) which consists of trains of several hundred pulses of an intensity of 100 to 400 Hz. We briefly introduce next some of the most common types of experiments in which HFS is used to induce LTP and investigate its relation to learning and memory.

One of the first experimental methodologies used to study how LTP might be involved in memory and learning consists in recording ESPSPs from learning animals and tracking whether the response changes after the learning experience. On the proposed classification of mechanistic experiments, this type of manipulation counts as a top-down experiment. The intervention applied on this experimental system is at the macrolevel of the rat's behaviour of exploring a new environment (traditionally a Morris water maze²¹), while the primary detection device targets the microlevel of the electrophysiological response of the neural cells in the dentate gyrus or CA1 region of the rat's brain.

It is worth noting that the experimental protocol *per se* does not rule out the pos-

¹⁹This reassertion of the methodological stance is in line with our neutrality on the question of whether constitution and causal dependence relations are metaphysically distinct or not.

²⁰It is also possible to measure population spikes which correspond to summed action potentials of neurons rather than membrane potentials. Since the experiments cited below do not apply this technique, we will not analyze it in more detail. For further references, see Holscher [2001]

²¹For more details on the battery of experimental protocols involving the Morris water maze, see Vorhees and Williams [2006] and Sullivan [2010].

sibility that the macrolevel intervention is a common cause for both the behavioural phenomenon of learning (a spatial task) and for the recorded LTP response. This implies that the correlation established via this kind of top-down experiment can be equally well accommodated by a common-cause and a constitutive model, as suggested by Baumgartner and Casini [2017]. The choice between the two models is underdetermined by the available evidence. The abductive view predicts that (1) only further experimentation which would systematically fail to de-couple LTP and behavioural phenomena can guarantee a constitutive inference showing that LTP is a component part of the molecular mechanism for memory; while (2) manipulations which de-couple the two levels would support the inference that LTP is not a component (constitutive part) of a memory mechanism.

Are these predictions confirmed by contemporary neurobiological experimental practices? On the one hand, critics of the synaptic mechanisms theory of memory (Shors and Matzel [1997]; Gallistel and Matzel [2013]; Gallistel and Balsam [2014]; Trettenbrein [2016]) point out that evidence from top-down behavioural experiments does not support the hypothesis that induction of LTP is either a necessary or sufficient condition for the storage of memories. On the other hand, there are still strands of LTP research that continue to investigate the link between microscopic LTP and cellular information storage as a form of memory. These experiments can be interpreted in line with the abductive account. Namely, as systematic attempts to break the common coupling between the macrolevel behavioural phenomena (for instance, spatial learning, navigation in specific types of mazes, conditioned suppression of activity, formation of spatial or olfactory memories) and microlevel elements of the postsynaptic mechanism for LTP (such as the role of CaM kinase or of tyrosine kinase in hippocampal LTP).

As an example from the latter category, we draw attention to experiments that explicitly investigate correlations between modulators of hippocampal LTP and memory (Shors and Matzel [1997]). What is interesting about these experiments is that they start off by acknowledging the common cause coupling between the two levels. For instance, chronic lead and alcohol consumption are known to impair memory storage and inhibit the induction of LTP. But inferring from behavioural manipulations that these compounds are detrimental to learning or memory as a result of their effect on hippocampal LTP is unwarranted. The hypothesis is underdetermined by the existence of other well-established mechanisms via which these substances influence memory storage capacities. For example, alcohol causes membrane fluidization and depresses both excitatory and inhibitory synaptic activity which eventually leads to disruption of normal function in the cerebral cortex. All or any of these effects could disrupt the information processing necessary for memory formation. Establishing that alcohol disruption of LTP blocks memory formation requires showing that the correlation between the two cannot be broken while varying all the other common cause processes via which alcohol might disrupt learning

and memory formation.

The correlation between LTP and higher level behavioural phenomena has also been studied using non-spatial learning paradigms, such as olfactory discrimination. For the rat, olfaction is a primary sensory modality, with olfactory information being partly processed by the hippocampal formation. Several types of experimental manipulations have revealed stable correlations between hippocampal LTP and olfactory learning (Shors and Matzel [1997]; Holscher [2001]). For example, using tetanic stimulation of the lateral olfactory tract as a discriminative cue, evoked responses in the piriform cortex were potentiated in the animals that learned the target discrimination.²² Furthermore, the cellular activity recorded in the same region during learning had a similar pattern to the tetanic stimulation used for the discrimination. Still, these experiments fail to be conclusive with regards to the constitutive relevance of LTP for memory formation. First, the correlations have been shown to be stable only for the olfactory system, so it is not likely that additional experimentation can establish that hippocampal LTP is necessary for memory processes in general. Second, the type of stimulation used in these experiments has been shown to induce LTP in the hippocampus even in animals that don't learn the target behaviour. According to the abductive account, the latter type of results challenge the conclusion that hippocampal LTP is constitutively relevant for memory formation because they seem to break the common cause coupling between the behavioural and the cellular level.

The top-down experiments sketched so far fail to establish conclusively the necessity of LTP for memory formation or learning. Even when stable correlations between hippocampal LTP and behaviour are discovered, experimental investigations fail to break the common cause coupling of the behavioural and cellular levels. The relatively small number of behavioural manipulations investigating the link between memory and LTP does not entail that it is a well-established fact that LTP is 'the memory mechanism', and that further experimentation is not required. As we suggested above, growing skepticism in the scientific community concerning the link between LTP and memory formation or learning seems to be based on experiments that decouple the two. However, one need not endorse the demise of the LTP hypothesis for memory to show that the abductive account provides a plausible reconstruction of top-down experiments in LTP and memory research. Unlike the MM account, the abductive account accommodates a wider variety of experimental practices, echoing the demand of some scientists for more behavioural manipulations investigating purported correlations between the macrolevel behaviours of memory formation and learning and the cellular and molecular mechanisms of LTP. The account also predicts that 'crucial' experiments, if any, will be considered those that manage to break the common cause coupling thus showing that LTP is not a constituent

²²Animals that were exposed to the simulation in a behaviourally irrelevant task did not learn nor did exhibit LTP.

of memory.

Turning to the category of bottom-up mechanistic experiments, we can point to several well-established experimental protocols. In behavioural pharmacological studies, researchers typically seek to establish correlations between the effect of drugs that inhibit LTP formation and the effect of these drugs on learning. An often cited study (Morris *et al.* [1986]) investigated the effect of AP5, a drug that blocks the NMDA glutamate receptor, on LTP and learning abilities. Doses of AP5 that blocked LTP induction *in vitro* and *in vivo* also seemed to impair learning of a spatial water maze task (Morris [1989]). However, additional experimentation could not validate the inference from this stable correlation to the necessity of LTP for memory formation. Subsequent studies showed that if the rats are pretrained in a spatial task in the water maze without drug application and then trained in a spatial task in another room, AP5 does not block learning even though it blocks the induction of LTP (Bannerman [1995]). A similar study showed that non-spatial pretraining in the water maze is sufficient to prevent the amnesic effect of NMDA antagonists in a subsequent water maze task (Saucier and Cain [1995]). These studies might be interpreted as showing that the drug induces motor impairments that the animal has to cope with before being able to learn the task. Therefore, the correlations established via this experimental methodology cannot be seen as providing direct evidence for LTP-like effects during learning. The direct inference is blocked by the fact that disturbing LTP and blocking learning with the same drug does not necessarily mean that both effects are based on the same mechanism.

Still one might argue that these bottom-up experiments support a different kind of inference which pertains to the dependence relation obtaining between LTP and learning. If appropriate experiments could show that animals learn while LTP induction in the hippocampus is blocked by a drug like AP5, then they would provide direct evidence for the hypothesis that the LTP mechanism is not necessary for learning. Otherwise put, in line with the abductive account, such experiments would establish that there is no common-cause coupling between macroscopic LTP and learning. However, they would not suffice to infer that the NMDA glutamate receptor (the target of the AP5 drug) is not constitutively relevant for the LTP response. Or that a different LTP mechanism is not inextricably coupled to learning phenomena. More systematic experimentation would be required to conclude that there are good empirical grounds for ruling out either the constitutive relevance of NMDA from a LTP mechanism or the constitutive dependence between LTP and learning. Even in that case, the evidential support would be at best indirect. That is, the results of behavioural pharmacological studies seem to be in line with the predictions of the abductive account in the sense that they underdetermine both types of constitutive inferences.

Another type of bottom-up experiment involves the saturation of all synapses in the hippocampus by electrical stimulation. If LTP is the mechanism supporting the

macrolevel behavioural phenomenon, then this experimental intervention would trigger the upregulation of all synapses which in turn should make learning impossible. Some of the first studies using this experimental methodology seemed to confirm this hypothesis, showing that after saturation of synapses animals were not able to learn a spatial task normally, while learning in animals in which LTP was induced but allowed to decay was not impaired (Castro *et al.* [1989]). While these initial experiments were taken to provide evidence in support of the stability of the correlation between LTP and learning, follow-up studies reported the opposite result: animals that showed more LTP after HFS learned the task even better (Jeffery and Morris [1993]). These latter results led to further refinements of the experimental procedure in order to make sure that synaptic saturation was indeed achieved, but the forthcoming results failed again to be conclusive (Holscher [2001]). In other words, these bottom-up experiments could not support constitutive inferences. On the contrary, as they seem to break the common cause coupling, they provide preliminary evidence for the constitutive irrelevance of LTP in learning tasks.

A different experimental methodology has been developed to check whether LTP can be elicited in the hippocampus at all. Experiments with gene deletion (knock-out) or transgenic mice strains use mutant mice strains that do not express a defined gene for a protein that has been previously shown to be crucial for LTP induction to compare their ability to learn spatial or non-spatial tasks. As in the case of behavioural pharmacological studies, the inferences drawn from these experimental methods are not conclusive. Even when such bottom-up experiments can establish a stable correlation between learning impairments and reduced or blocked LTP formation, they fall short of providing direct evidence that the same mechanism underlies both the blocking of LTP and learning impairments. Knocking out a gene that is essential for neuronal activity might impair LTP without also being involved in learning impairments. That is, the evidence that these experiments makes available underdetermines the choice between LTP and non-LTP mechanistic models of learning. Moreover, while several experiments on mutant mice strains that don't express particular genes showed strong correlations between LTP development and learning impairments, others failed to replicate the result and showed that mice in which LTP was not inducible were still able to learn a spatial task in the water maze (Malenka *et al.* [1989]; Grant *et al.* [1992]; Huang *et al.* [1995]).

In the attempt to improve existing types of bottom-up experimental studies, neuroscientists worked with experimental systems which involved selective knock-outs (of the expression of NMDA receptors) in a limited brain area (CA1) of transgenic mice. These surgical knock-out techniques were meant to avoid unwanted indirect effects caused by lack of NMDA receptors in other parts of the brain or by other developmental impairments in the brain. The experimental studies reported that mice were indeed impaired in learning spatial tasks and that LTP of field EPSPs could not be induced in these animals (Wilson and Tonegawa [1997]). These studies provide some of the most stable

correlations between learning abilities in animals and LTP inducibility. However, they do not yet establish the constitutive relevance of LTP for macrolevel phenomena of learning and memory. Neuroscientists continue to explore possible ways in which the cellular and molecular levels could be experimentally manipulated to bring about behaviourally observational effects. These efforts could be interpreted along the lines of the abductive account. If LTP is constitutively relevant for memory formation or learning, then new experimental manipulations should consistently fail to decouple the two levels. Revisions of the hypothesis concerning the involvement of LTP in memory and learning processes, on the other hand, will follow from experiments which would succeed in breaking the common cause coupling of the two levels.

5.2 A comparative summary

The examples reviewed above illustrate the methodology of interlevel mechanistic experiments in virtue of the fact that they explore the link between the macrolevel behavioural phenomena of memory formation or learning and microlevel LTP phenomenon. We claim that the abductive account provides a plausible reconstruction of the main trends in this area of neurobiological research. On the one hand, the view provides an interpretation for the rising skepticism against the LTP-based hypothesis of memory. Interlevel experiments that succeed to break the common cause coupling between the two levels support the idea that LTP might not be constitutively relevant for memory or for learning. On the other hand, the view also makes sense of the call for more experimentation that would make constitutive inferences (or the inferences about the constitutive irrelevance of LTP mechanisms to memory) more robust.

In addition, the abductive account sharpens the methodological distinction between interlevel constitutive experiments and causal experiments by highlighting the special type of underdetermination problems facing constitutive inferences. According to Baumgartner and Casini [2017], the evidence produced by interlevel experiments cannot empirically distinguish between a pure common cause model and a constitutive model of the relationship between LTP and memory. The choice between the two types of models is underdetermined by anything that falls short of a complete exploration of the whole space of possible ways of breaking the coupling between the two levels. However, the correlations revealed by interlevel experiments can be taken to support indirectly a constitutive model. While the pure common cause model does not provide any insight into why interlevel correlations cannot be broken, the constitutive model postulates that this sort of stability stems from a particular type of dependence that obtains between the two mechanistic levels under investigation. In other words, failure to de-couple the levels counts indirectly as evidence for a constitutive dependence relation linking them. Thus, inferring abductively a constitutive dependence relation is the only way of blocking the

underdetermination problem for constitutive inferences.

The latter consequence gives the abductive account one important advantage over Craver's own reconstruction of constitutive experiments in neuroscience (Craver [2002], [2007a], [2007b]). Although the MM account can accommodate the call for more interlevel experimentation in terms of neuroscientists' concern for the robustness of their experimental results, the abductive view does a better job at capturing the differences between how scientists seem to treat interlevel or constitutive inferences and causal inferences. More specifically, the latter account traces the hesitancy to fully endorse constitutive models linking LTP and memory to the special type of underdetermination problem facing constitutive models in general. In contrast, causal models of the molecular mechanisms of LTP itself are more fully endorsed by the neuroscientific community even when underdetermination problems are still to be encountered (cf. Sanes and Lichtman [1999]; Granger and Nicoll [2014])

In relation to this point, it is worth noting that for the aim of promoting LTP and memory experimental research as being one of the best case studies for illustrating the methodology of interlevel experiments and its links to constitutive reasoning in the biological sciences, Craver focuses too much too often on causal relevance experiments instead. In fact, his analysis targets frequently the description of the causal aspects of the microscopic LTP mechanism (Craver [2002], [2007b]; Craver and Darden [2013]; see also Fagan [2015]). The conclusive character of the inferences concerning the causal relevance of different mechanistic components of the LTP mechanism is however not transferable to inferences about the relation between the microlevel of the LTP mechanism and the macrolevel of memory and learning behaviours. Indeed Craver shies away from claiming that scientists have developed any crucial experiment to establish that LTP is a constitutively relevant component of memory mechanisms. Instead he seems to take the practical limitations of interlevel experiments to block such a definitive conclusion. Although descriptively accurate with respect to the current experimental landscape of LTP and memory research, this stance is also at odds with the implication of the MM approach that constitutive relations between higher and lower-level mechanistic levels can be established conclusively on the basis of the cumulative evidence of pairs of top-down and bottom-up experiments. We thus propose that Craver's approach is less convincing than the abductive account according to which the evidence yielded by interlevel experiments will always underdetermine constitutive inferences.

The abductive account is not the only attempt to improve the philosophical analysis of constitutive experimentation and reasoning in science. For instance, Harbecke [2015b] articulates a regularity theory of constitution and constitutive inference, while Gebharter [2017] proposes a modified version of a causal search algorithm for detecting constitutive relations that are formally similar (under certain conditions) to direct causal relations. The abductive account shares with Harbecke [2015b] the insight concerning the induc-

tive risks implied by constitutive inferences. Although the abductive account offers a more coarse grained assessment of the structure of interlevel experiments in neuroscience than the one entailed by Harbecke’s set of inference rules, it also avoids the problems confronting regularity views of causation more generally. Gebharter’s [2017] account on the other hand does not rely either on strict regularities nor on interventions to analyze constitutive inferences. In order to establish whether such an account entails a more informative reconstruction of neurobiological interlevel experiments, one must show that the dependence relations they target satisfy the formal conditions of the causal search algorithm used to model constitutive inferences. It will be interesting to show in future research which of the three accounts provides more accurate analyses and predictions about the experimental practices of neurobiology. The aim of this paper was to show that the abductive account provides a plausible reconstruction of interlevel experiments and constitutive reasoning in at least one important area of neurobiological research.

6 Conclusions

On the abductive theory of constitution, interlevel mechanistic experiments support constitutive inferences and explanations only by providing indirect evidence about constitutive relevance relations. Interlevel experiments fail to de-couple the common causes for the mechanistic macro and micro-levels. Positing a non-causal dependence relation between the macrolevel and the microlevel, that is a constitutive relevance relation, aids in explaining this systematic de-coupling failure. Thus, constitutional models which postulate such dependence relations are explanatorily superior to pure common cause models which can be built on the back of the same types of interlevel experiments. In other words, interlevel experiments cannot provide direct evidence that would distinguish empirically a pure common cause model from a constitutive model of a mechanistic phenomenon. However, a large class of interlevel experiments (top-down, bottom-up and mixed experiments) yields indirect evidence that supports the explanatory value of constitutive models in neuroscience. This characterization clearly distinguishes constitutive (relevance) inferences from causal (relevance) inferences. In the case of the latter, it is possible (at least in principle) to devise an experiment that would undercut the underdetermination of the causal inference, whereas constitutive inferences suffer from a more radical form of underdetermination. This is due to the fact that it is practically unfeasible to explore the whole space of logically possible interlevel experiments.

So, instead of claiming, with Craver, that pairs of top-down and bottom-up mechanistic experiments can conclusively establish mechanistic constitution relations, Baumgartner and Casini’s recommendation is to view the evidence provided by these experiments as offering indirect or abductive support for a constitutive claim linking different mechanistic levels. In the last section, we argued that this recommendation is in line with

the experimental methodologies observed in at least one important area of neuroscientific research.

The abductive view of constitutive experiments has also two broader implications for the philosophy of experimentation (Weber [2014]; Franklin and Perovic [2016]). The first one concerns the types of evidence that mechanistic experiments can be said to provide in support of different research hypotheses. The abductive account highlights the fact that not all experiments yield direct evidence for a particular hypothesis or model under investigation. The fact that some inferences are bound to be underdetermined by the data generated through some types of experiments does not make the latter useless. Rather, according to the abductive account, these experiments generate indirect evidence for certain types of scientific inferences. Secondly, the abductive account suggests that similar types of experiments (causal and constitutive relevance experiments) can be used to support and refine different types of scientific inferences. This calls for further investigations into the different inferential roles of experiments in scientific investigation.

Works Cited

- [1] Anderson, P., Morris, R., Amaral, D., Bliss, T. and O'Keefe, J. [2006]: *The Hippocampus Book*, New York: Oxford University Press.
- [2] Bannerman, D.M., Good M. A., Butcher S.P., and Ramsay M. [1995]: 'Distinct components of spatial learning revealed by prior training and NMDA receptor blockade', *Nature*, **378**, pp. 182-6.
- [3] Baumgartner, M. and Casini, L. [2017]: 'An abductive theory of constitution', *Philosophy of Science*, **84.2**, pp: 214-33.
- [4] Baumgartner, M. and Gebharder, A. [2016]: 'Constitutive Relevance, Mutual Manipulability, and Fat-Handedness', *British Journal for the Philosophy of Science*, **67.3**, pp. 731-56.
- [5] Bechtel, W. [2017]: 'Explicating top-down causation using networks and dynamics'. *Philosophy of Science*, **84.2**, pp. 253-74.
- [6] Craver, C. and Bechtel, W. [2007]: 'Top-down causation without top-down causes', *Biology and Philosophy*, **22.4**, pp. 547-63.
- [7] Craver, C.F. [2002]: 'Interlevel experiments and multilevel mechanisms in the neuroscience of memory', *Philosophy of Science*, **69.S3**, pp. 83-97.
- [8] Craver, C.F. [2003]: 'The Making of a Memory Mechanism', *Journal of the History of Biology* **36**, pp. 153-95.

- [9] Craver, C. F. [2007a]: ‘Constitutive explanatory relevance’, *Journal of Philosophical Research*, **32**, pp. 3-20.
- [10] Craver, C. F. [2007b]: *Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience*, New York: Oxford University Press.
- [11] Craver, C. F. and Darden, L. [2013]: *In Search of Mechanisms: Discovery across the Life Sciences*, Chicago: University of Chicago Press.
- [12] Dudai, Y. [1995]: ‘On the relevance of long-term potentiation to learning and memory’. In Weinberger, N.M., McGaugh, J.L. and Lynch, G. (eds.) *Brain and Memory: Modulation and Mediation of Neuroplasticity*. New York: Oxford University Press.
- [13] Fagan, M.B. [2013]: *Philosophy of Stem Cell Biology. Knowledge in Flesh and Blood*. New York: Palgrave Macmillan.
- [14] Fagan, M. B. [2015]: ‘Collaborative explanation and biological mechanisms’, *Studies in History and Philosophy of Science Part A*, **52**, pp. 67-78.
- [15] Franklin, A. and S. Perovic. [2016]: ‘Experiment in Physics’. In E. Zalta (ed.) *Stanford Encyclopedia of Philosophy*, URL = <https://plato.stanford.edu/archives/win2016/entries/physics-experiment/>
- [16] Gallistel, C.R. [1995]: ‘Is long-term potentiation a plausible basis for memory?’ In Weinberger, N.M., McGaugh, J.L. and Lynch, G. (eds.) *Brain and Memory: Modulation and Mediation of Neuroplasticity*. New York: Oxford University Press.
- [17] Gallistel, C.R. and Balsam, P.D. [2014]: ‘Time to rethink the neural mechanisms of learning and memory’, *Neurobiology of Learning and Memory*, **108**, pp. 136-44.
- [18] Gallistel, C.R. and Matzel, L.D. [2013]: ‘The Neuroscience of Learning: Beyond the Hebbian Synapse’, *Annual Review of Psychology*, **64**, pp: 169-200.
- [19] Gebharter, A. [2017]: ‘Uncovering Constitutive Relevance Relations in Mechanisms’, *Philosophical Studies* **174.11**, pp. 2645-66.
- [20] Granger, A. J. and Nicoll, R.A. [2013]: ‘Expression Mechanisms Underlying Long-Term Potentiation: A Post-synaptic View, 10 years on’, *Philosophical Transactions of the Royal Society B*, **369**, pp. 20130-6.
- [21] Harbecke, J. [2010]: ‘Mechanistic Constitution in Neurobiological Explanations’, *International Studies in the Philosophy of Science*, **24**, pp: 267-85.
- [22] Harbecke, J. [2014]: ‘The role of supervenience and constitution in neuroscientific research’, *Synthese*, **191**, pp. 725-43.

- [23] Harbecke, J. [2015a]: ‘Regularity Constitution and the Location of Mechanistic Levels’, *Foundations of Science*, **20.3**, pp: 323-38.
- [24] Harbecke, J. [2015b]: ‘The Regularity Theory of Mechanistic Constitution and a Methodology for Constitutive Inference’, *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, **54**, pp: 10-9.
- [25] Holscher, C. [2001]: ‘Long-term potentiation as a model for memory mechanisms: The story so far’. In Holscher C. (ed.) *Neuronal Mechanisms of Memory Formation*. New York: Cambridge University Press.
- [26] Huang, Y.Y., *et al.*[1995]: ‘A genetic test of the effects of mutations in PKA on mossy fiber LTP and its relation to spatial and contextual learning’. *Cell*, **83**, pp. 1211-22.
- [27] Illari, P. and Willimson, J. [2012]: ‘What is a mechanism? Thinking about mechanisms across the sciences’, *European Journal for Philosophy of Science*, **2.1**, pp. 119-35.
- [28] Jeffery, K.L. and Morris, R.G.M. [1993]: ‘Cumulative long-term potentiation in the rat dentate gyrus correlates with, but does not modify, performance in the water maze’, *Hippocampus*, **3**, pp. 133-40.
- [29] Kaplan, D.M. [2012]: ‘How to Demarcate the Boundaries of Cognition’, *Biology and Philosophy*, **27**, pp. 545-70.
- [30] Laudan, L. and Leplin, J. [1991]: ‘Empirical Equivalence and Underdetermination’, *The Journal of Philosophy*, **88.9**, pp. 449-72.
- [31] Leuridan, B. [2012]: ‘Three Problems for the Mutual Manipulability Account of Constitutive Relevance in Mechanisms’, *British Journal for the Philosophy of Science*, **63**, pp. 399-427.
- [32] Levy, A. [2016]: ‘The Unity of Neuroscience: A Flat View’, *Synthese*, **193.12**, pp. 3843-63.
- [33] Malenka, R.C., *et al.* [1989]: ‘The impact of postsynaptic Ca on synaptic transmission and its role in LTP’, *Trends in Neuroscience*, **12**, pp. 444-50.
- [34] Morris, R.G.M. [1989]: ‘Synaptic plasticity and learning: Selective impairment of learning in the rat’, *Journal of Neuroscience Methodology*, **11**, pp. 47-60.
- [35] Morris, R.G.M., *et al.* ‘Selective impairment of learning and blockade of long-term potentiation in vivo by NMDA receptor antagonist AP5’, *Nature*, **319**, pp. 774-6.

- [36] Okasha, S. [2002]: ‘Underdetermination, holism and the theory/data distinction’, *Philosophical Quarterly*, **52.208**, pp. 303-19.
- [37] Sanes, J.R. and Lichtman, J. W. [1999]: ‘Can Molecules Explain Long-Term Potentiation?’, *Nature Neuroscience*, **2.7**, pp. 597-604.
- [38] Saucier, D. and Cain, D.P. [1995]: ‘Spatial learning without NMDA receptor-dependent long-term potentiation’, *Nature*, **378**, pp. 186-9.
- [39] Schindler, S. [2013]: ‘Mechanistic Explanation: Asymmetry Lost’. In Dieks, D. and Karakostas, V. (eds.) *Recent Progress in Philosophy of Science: Perspectives and Foundational Problems*. Springer.
- [40] Scholl, R. and Raz, T. [2016]: ‘Towards a Methodology for Integrated History and Philosophy of Science’. In Sauer, T. and Scholl, R. (eds.) *The Philosophy of Historical Case Studies. Vol. Boston Studies in the Philosophy and History of Science*. Springer.
- [41] Schors, T. and Matzel, L.D. [1997]: ‘Long-term Potentiation: What’s Learning Got to Do With It’, *Behavioral and Brain Sciences*, **20**, pp. 597-655.
- [42] Sullivan, J. [2009]: ‘The multiplicity of experimental protocols: A challenge to reductionist and non- reductionist models of the unity of neuroscience’, *Synthese*, **167**, pp. 511-39.
- [43] Sullivan, J. ‘Reconsidering “Spatial Memory” and the Morris Water Maze’, *Synthese*, **177**, pp. 261-83.
- [44] Trettenbrein, P. C. [2016]: ‘The Demise of the Synapse As the Locus of Memory: A Looming Paradigm Shift?’, *Frontiers in Systems Neuroscience*, **10**, p.88.
- [45] Vorhees, C. V. and Williams, M.T. [2006]: ‘Morris Water Maze: Procedures for Assessing Spatial and Related Forms of Learning and Memory’, *Nature Protocols*, **1.2**, pp. 848-58.
- [46] Weber, M. [2014]: ‘Experiment in Biology’. In E. Zalta (ed.) *Stanford Encyclopedia of Philosophy*, URL: <http://plato.stanford.edu/archives/win2014/entries/biology-experiment/>.
- [47] Werndl, C. [2013]: ‘On choosing between deterministic and indeterministic models: underdetermination and indirect evidence’, *Synthese*, **190.12**, pp. 2243-65.
- [48] Wilson, M.A. and Tonegawa, S. [1997]: ‘Synaptic plasticity, place cells and spatial memory: Study with second generation knockouts’, *Trends in Neuroscience*, **20**, pp. 102-5.

- [49] Wimsatt, W. [1994]: ‘Robustness, Reliability, and Overdetermination’. In Brewer, M. and Collins, B. (eds.) *Scientific Inquiry and the Social Sciences*. San Francisco: Jossey-Bass, pp. 259-91.
- [50] Woodward, J. [2003]: *Making Things Happen: A Theory of Causal Explanation*. New York: Oxford University Press.
- [51] Woodward, J. [2015]: ‘Interventionism and Causal Exclusion’, *Philosophy and Phenomenological Research*, **91.2**, pp. 303-47.
- [52] Ylikoski, P.[2013]: ‘Causal and Constitutive Explanations Compared’, *Erkenntnis*, **78.2**, pp. 277-97.