PAIN IS NOT A NATURAL KIND

by

JENNIFER CORNS

A dissertation submitted to the Graduate Faculty in Philosophy in partial fulfillment of the requirements for the degree of Doctor of Philosophy, The City University of New York

2012

This manuscript has been read and accepted for the Graduate Faculty in Philosophy in satisfaction of the dissertation requirement for the degree of Doctor of Philosophy.

_____	_____
Date	David Rosenthal Chair of Examining Committee
_____	_____
Date	John Greenwood Executive Officer

Jesse Prinz, Advisor

Michael Devitt

Peter Godfrey-Smith

Stephen Neale

THE CITY UNIVERSITY OF NEW YORK

Abstract

PAIN IS NOT A NATURAL KIND

by

Jennifer Corns

Advisor: Jesse Prinz

Pain is central to our lives. Despite that, I argue, pain is not a natural kind. Chapter 1 identifies a kind as natural insofar as it is usefully referred to in the generalizations of the relevant science(s). Following Boyd, I take the best indicator of this to be causally interlinked clusters of properties, and I update Boyd's approach by relating it to recent work on mechanistic explanation. Chapter 2 employs clinical observations to argue that philosophical, unitary accounts of pain are inadequate. Pain is a multidimensional experience that paradigmatically includes sensation, perception, emotion, cognition, and motivational responses. This multidimensionality is included in the dominant scientific models of pain to which I turn in Chapter 3, but none identify a mechanism (neurobiological or otherwise) underlying the usual co-occurrence of pain's clustering properties. Chapter 4 argues that recent research attempting mechanism-based classifications of pain types allows us to conclude that each token pain is determined by an *idiosyncratic* convergence of the activity of *multiple* mechanisms. Neither pain nor any type of pain is a natural kind.

Once it is established that a folk kind is not natural, it is customary to become an eliminativist or a pluralist; I resist both options and offer an alternative. Chapter 5 characterizes pluralism as the position that the mechanical heterogeneity of a kind does not undermine its naturalness—either because there is homogeneity at another level, or because the heterogeneity is negligible relative to the target phenomenon. I argue against both options. In the final chapter, I argue that pains are

very real despite their non-naturalness; eliminativism should be resisted. The idiosyncrasy argued to subvert scientific generalizations mentioning pain does not disrupt utile reference to pain in everyday life. Neither does idiosyncrasy entail non-existence. The dissertation as a whole may then be considered a case study of a robust and important folk-psychological kind that scientific inquiry reveals is not natural.

Acknowledgements

There are too many people to thank. These include:

My family: my mother—who first taught me to insatiably ask questions, my father—who first taught me to analyze and apply texts, and my sister—who first hit me in the face for talking too much. Numerous mentors: Michael Matriotti—who introduced me to the study of philosophy, Phillip Goggans—who convinced me there might be a living in this racket, Kenneth Einar Himma—who has continually, incredibly, believed in and supported me, Stefan Baumrin—who understood me when I felt a stranger in a strange world, David M. Rosenthal—who introduced me to the wonders of the mind and how to begin plumbing their depths, and Michael Devitt—who gave me many helpful comments and showed by example that it's OK to be a philosopher who laughs. My dissertation committee—for asking tough questions and pointing out holes. My friends: David Pereplyotchik—who regularly listened to me rant and cry when we should have been sleeping, Monique Whitaker—who taught me that it's alright to be human, and Charlotte Latham—who is the guardian of my sanity.

My special thanks go to the Kanet family, for lending me their beach house, where most of this work was written.

Last and most important, of course, many thanks belong to my advisor Jesse Prinz. Jesse's advice, support, regular feedback, comments, and insight, were necessary for this piece to be possible.

Table of Contents

List of Figures

Chapter 1: Natural Kinds and Mechanisms

1. Introduction

Pain is central to our lives. We go to great lengths to avoid it and would be in serious trouble without it. Pain can serve as an important warning. When in pain it can be difficult, sometimes impossible, to attend to anything else. Despite the presence and importance of pain in our lives, the central thesis of this dissertation is that pain is not a natural kind.

The first task in establishing this thesis is to develop an account of natural kinds. A quick survey of the literature reveals little agreement. Even basic questions—what natural kinds *are* and how to substantiate whether a kind is natural, or a term is a natural-kind term—are disputed. If we are to care whether pain, or anything else, is a natural kind, we need answers to these questions.

The most obvious starting point is to hold that natural kinds are those usefully referred to in the relevant sciences' generalizations.[1] I argue in this chapter that, though requiring further refinement, this is the *only* appropriate criterion for a kind to be natural. Moreover, any term used to refer to such a kind is a natural-kind term. Once it is been determined that a kind and its term are earning, or are credibly expected to earn, their scientific keep, further disputes about whether the kind in question deserves the honorific title "natural" are uninteresting. The interesting issue is whether a kind (whatever we call it) is useful for scientific theorizing. I will dub this criterion the *scientific utility criterion*.[2]

The scientific utility criterion might appear to leave out the traditional requirement that natural kinds have specified, or specifiable, essences—whether non-natural kinds also have

---

[1] Being referred to in generalizations is meant to rule out things like beakers and telescopes, which presumably form a kind useful for scientific inquiry that are not themselves its subject.

[2] Thanks to Michael Levin for this apt title.

essences or not. Seeming essentialists currently include such important figures as Putnam (1975) and Kripke (1980). In section 2, I examine this strand of the natural-kinds tradition and argue that if essences are good for identifying natural kinds, it is because they are good for identifying those kinds usefully referred to in science's generalizations. Emphasizing whether a term and the kind it refers to are useful for doing science will, with respect to the purported kind, deliver all that focusing on essences was properly employed to do within the natural kinds tradition.

Using the scientific utility criterion to determine naturalness requires saying more, however, about when a kind and the term referring to it are useful for scientific inquiry. Sections 3, 4, and 5 provide the requisite elaboration that I will employ for the remainder of the dissertation. I turn first to Quine, concluding in section 3 that a kind's naturalness is a matter of degree, relative to the success of the generalizations that reference to the kind facilitates. When are these successful? To answer this question, I turn in section 4 to Ian Hacking's (1991) identification of 4 conceptions of natural kinds, all of which he thinks should be embraced: Russell-kinds, Mill-kinds, Pierce-kinds and Leibniz-kinds. Each conception, I argue, is best understood as a supplement to its predecessor, and allows for increasingly successful generalizations. The as yet final supplement in this tradition is Richard Boyd's influential conception of natural kinds as those exhibiting a cluster of properties whose usual co-occurrence is explained by an underlying, homeostatic mechanism. Section 5 explains and elaborates these kinds and, in keeping with Hacking's terminology, I will refer to them as Boyd-kinds. While there are no necessary and sufficient conditions for natural kindhood, I conclude that Boyd-kinds provide general indicators for determining whether a kind and its term are useful for scientific theorizing.

Boyd-kinds posit an underlying homeostatic mechanism, so in section 6 I focus on mechanisms and mechanistic explanation. I briefly consider 3 influential models of mechanistic explanation recently advanced, respectively, by Bechtel, Craver, and Glennan. Here I will conclude that a mechanism can be understood liberally as a whole with *parts* that are *organized* and *do things*, such that the organization and goings-on of the parts determines some behavior or phenomenon. A mechanistic explanation explains a phenomenon by positing a mechanism that determines it.

Sections 7 and 8 close the chapter and prepare the ground for those following. In section 7, I briefly discuss the recent movement in pain science towards a mechanism-based classification of pain. In light of the account of natural kinds and mechanistic explanation developed here, it is apparent that this recent movement is an attempt to identify many pain *types* as natural kinds. Two of the problems Craver identifies for developing mechanistic explanations are under-specification and lumping. In section 8, I argue that these problems are evident in the motivation for mechanism-based classifications of different types of pain, and the major problem facing their provision. This exposition reveals two distinct claims concerning pain's non-naturalness: The different types of pain do not jointly constitute a natural kind and no type of pain is itself a natural kind. Chapters 2, 3, and 4 argue for these claims.

2. Of Essences

Essences are understood in multiple ways. They may be taken to be a thing's definition, necessary and sufficient conditions, micro-structural properties, intrinsic properties, or some combination thereof. I will not take a position on which, if any, of these conceptions of essences

is correct. Instead, in this section I will argue only for the following claim: If the requirement that *natural kinds* have specifiable essences or that members belong to their natural kinds essentially is ever appropriate, it is because essence-identification facilitates classifications useful for scientific theorizing. In turning to the tradition, we will see that there is widespread agreement about this. Scientific utility is identified by the essence theorists themselves as one, if not *the,* central purpose of essence-identification. It is irrelevant here whether essences are useful in other inquiries for other reasons. Having established that this is the only legitimate role for essences *relevant to natural-kind questions*, we can focus on a kind's scientific utility without settling whether the kind or its members have specified, or specifiable, essences. For natural-kind determination, there is no need to talk further about essences.

To establish the relationship between essences and the scientific utility criterion I prefer, I focus on an ancient, modern, and contemporary essence theorist: Aristotle, Locke, and Devitt. Despite the variability in their accounts of what essences *are*, all identify the *role* of essences by their scientific utility.

For Aristotle, substances and, more contentiously, nonsubstantial particulars, were defined by their essences. Aristotelian essences are the "what it is" or "what it is to be" of a thing. We define a thing by specifying its essence.[3] The essence of each substance is correlative to its *species*, which is subsumed under a *genus*.[4] For Aristotle, the categorization of things by their essences provided groupings intended for use in research and inquiry.

It would be anachronistic to say that Aristotle took having a specifiable essence to be a necessary condition for a set to be a *natural kind*. It is commonplace, however, to understand

---

[3] For a modern incarnation of an approach that takes essences to be definitions of things, see Fine, "Essence and Modality: The Second Philosophical Perspectives Lecture," 1994.
[4] Aristotle, *Metaphysics*, Z.

him as using his hierarchy of essences as his categories for scientific theorizing. Aristotelian essences are most centrally employed for scientific theorizing.

Locke took objects to have both real and nominal essences. The real essence is the unobservable constitution of a thing responsible for its observable qualities. The nominal essence is a construction from our abstract ideas of it formed by observations of its qualities. According to Locke, we can only categorize according to the constructed ideas abstracted from our observations. These constructions constitute our artificial, but non-arbitrary, human categorizations of objects. Sorts and species are paradigmatic nominal essences, while real essences remain unknown.[5]

Locke's irritation with those, like Aristotle, who take species or sorts to specify real, as against merely nominal, essences is illuminating for our current inquiry. The Lockean objection is that doing so has "… very much perplexed the knowledge of natural things."[6] Three-legged dogs are not a problem unless we think it is part of the *real* essence of dogs to have four. If we take our categorizations to reveal real essences, however, any clash between the essence and the observable qualities will leave us wholly at a loss to reconcile the mismatch. We can observe only the number of legs any given dog has. Most dogs, as it turns out, have four, and we construct the nominal essence of the kind dog accordingly. Whether it is the *real* essence of a dog to have three, four, or any other number of legs cannot be settled by observation and, on Locke's view therefore, cannot be settled at all. In this way, real essences are "…*wholly useless* and unserviceable to any part of our knowledge…."[7] As such, we ought to turn instead to nominal essences and "… content ourselves with such *essences* of the sorts or species of things,

---

[5] Locke, *An Essay Concerning Human Understanding*, book 3, especially chapters 3 and 6.
[6] Ibid., 186.
[7] Ibid., 186.

as come within the reach of our knowledge."[8] While Locke believes that real essences must exist to determine the qualities that we can observe, it is our human categorizing of those observable qualities that allows our knowledge to advance.

Whether we should agree with Locke or not, I wish to stress his emphasis on what essences are *good for:* The advancement of our knowledge of nature. Real essences were, for Locke, "wholly useless." Why? Because they did not help advance our knowledge of the nature of things. Nominal essences, on the other hand, were useful because they did just that. While 'natural kind' is not a part of Locke's vocabulary any more than Aristotle's, the moral should be obvious. If we understand essences in a Lockean way, and require that natural kinds have specifiable essences, natural kinds are those determined by their nominal essences. Nominal essences, in turn, are an abstract categorization *useful for* advancing our knowledge of the nature of things—that is, useful for scientific inquiry.

Michael Devitt's endorsement of essences directly depends on their purported role in scientific inquiry. For him, an essence is whatever it is in virtue of which a thing is the kind of thing it is. These essences are known *a posteriori*. On Devitt's view, natural kinds have essential properties and "… it is a task of science to discover which these properties are."[9] Among other things, scientists are in the business of discovering essences, and their role in scientific theorizing is what renders them indispensable.

Though Devitt maintains that even non-natural kinds have essences, all of his arguments concerning essences are explicitly based on their role in scientific theorizing. His stated motivation for essentialism is the necessity of essences for explaining the observable properties

---

[8] Ibid., 186.
[9] Devitt, *Realism and Truth*, 22.

of groups of organisms.[10] He argues that biological explanations, structural explanations in particular, demand essences.[11] Furthermore, he claims that there must be some underlying intrinsic essence at least partly responsible for determining the truth of at least some generalizations.[12] His defense of the very existence of essences depends solely on their commensurability with contemporary scientific practice. Devitt allows that essences may be purely intrinsic, partly intrinsic, and partly extrinsic/relation, or purely extrinsic/relational—liberality he believes is necessary if essences are to be compatible with scientific theories. [13]

These three case studies from the history of philosophy help make clear that characterizing natural kinds by their scientific utility is at least *consistent* with those in the natural kind tradition that requires essences. Using scientific utility as the sole criterion helps, furthermore, to dispel the fear that accepting natural kinds commits one, via essences, to superstition or mystery. Insofar as the *natural kinds tradition* focuses on essences, it does so centrally to identify those kinds useful for scientific theorizing. Insofar as essence-identification does anything further, that furtherance is irrelevant to natural kindhood. The scientific utility criterion for natural kinds is in no tension with that strand of the tradition that emphasizes essences.

If we are not to look to essences, however, more needs to be said about when reference to a kind in a generalization of science is useful. Quine develops the idea that natural kinds are simply those that successfully support science's meta-inductions.[14] For explicating the

---

[10] Devitt, "Resurrecting Biological Essentialism," 365.
[11] Ibid., 349-355.
[12] Ibid., 352.
[13] Ibid., 345
[14] Quine's discussion owes a great deal to Goodman's. I have chosen to focus on the former since it is more focused on natural kinds in particular and less on predication or projectability in general. For further development of the notion of projection one should, of course, refer to

characterization of natural kinds by their scientific utility, it is a good place to start.

3. Quine

In "Natural Kinds," Quine sees the notions of kind and of similarity as intertwined, and attempts to define them clearly. A single similarity relation will not yield a definition of a kind, but comparative similarity will help; a kind has its members in virtue of their being *more* similar to each other than to any non-member. Quine concludes that "[i]f we reassess something *a* as less similar to *b* than to c, surely we will correspondingly permute *a, b, and c* in respect of their assignment to kinds; and conversely."[15] This definition is still imprecise and leaves comparative similarity undefined.

Despite this seeming inability to precisify, Quine maintains that we depend on our understanding of similarity, and therefore of kinds, for our everyday practices and all learning. In Quine's words, "… every reasonable expectation depends on resemblance of circumstances, together with our tendency to expect similar causes to have similar effects."[16] Whence the requisite resemblances? Enter innate similarity spaces. According to Quine, *without* innate similarity spaces, all stimuli would be "equally alike and equally different,"[17] but *with* them, we can develop habits and form expectations. All learning requires the formation of these habits and expectations, and induction just *is* the formation of these. On Quine's view, then, all learning is an implicit case of induction ultimately resting on unlearned classifications.

To illustrate: consider dogs. I learn what dogs are through exposure to dogs, during at

---

Goodman's *Fact, Fiction and Forecast*. My purpose here is merely to expand on when reference to a kind is useful for science, and Quine's discussion is better suited for that somewhat restricted purpose.

[15] Quine, "Natural Kinds,"161.

[16] Ibid., 157.

[17] Ibid., 162

least some of which an English speaker ostends an animal and says 'dog'. Also necessary is exposure to foils. I point, for instance, to a cat, say 'dog,' and am corrected. Given the potentially infinite number of properties a dog has in common with other things, however, I must have some *innate* grip on what the relevant properties for doghood might be. Without that, I would be unable to predict which things, in addition to the ostended member(s), are members of the kind I am learning.[18] After enough exposure to dogs and foils, I form habits and develop expectations about what *sort* of things dogs are. I have, that is, performed an implicit induction, and repeated inductions yield my understanding of the kind *dog*.

On Quine's view, however, there is no guarantee that the kinds resulting from these inductions are going to match any of the kinds *in nature*. To determine which kinds are natural we must *project,* and we must modify our understanding of the relevant similarity according to the success of those projections. We must, that is, theorize about which kinds exist in nature, test those theories, and revise them. This is taken to be a large part of what the scientist is doing: "By the trial-and-error process of theorizing he has regrouped things into new kinds which prove to lend themselves to many inductions better than the old."[19] Our notion of the kind is revised—i.e., our understanding of the relevant comparative similarity that delineates the kind is refined—in response to the success or failure of the meta-inductions of the scientists. We thought whales were fish, but upon the scientists' projections of the notions of fish and whale, we discovered that they are less similar than we thought. We therefore no longer project properties common among fish to whales. This allows for more successful inductions about both fish and whales.

---

[18] Note that on my reading of Quine's proposal it is *members* and not kinds that are ostended, though the *word* used in the ostension may be either a mass noun or a noun of divided reference—either of which may later be used by speakers to refer to kinds. This goes some way, it seems to me, towards assuaging the sort of worries presented by Fodor in "What is Universally Quantified And Necessary And A Posteriori And It Flies South In The Winter?"

[19] Ibid., 167.

Not all kinds used in scientific theorizing originate from implicit inductions based upon innate similarity spaces, nor does Quine require this. Some theoretical positing of kinds originates from scientific theory itself. Nonetheless, the account is unitary: Scientific inquiry performs inductions using notions of kinds, which are themselves outputs of previous inductions, and revises any first-order groupings in response to the success or failure of projections.

When meta-inductions are successful, the featured kinds are natural. Quine's kinds, however, may be more appropriately considered *more or less* natural. "Discovering" a kind is a matter of refining a previous notion of the relevant comparative similarity to improve the success of the subsequent meta-inductions.[20] Just how many successful second-order inductions a kind must feature in to qualify as *natural*, is a matter of degree and not (pardon the pun) of kind. Given that on Quine's view similarity is likewise always a matter of degree, this is to be expected.

Repeated meta-inductions and the refinements they demand allow scientists to arrive at refined accounts of the *mechanism* responsible for the unrefined notion of comparative similarities with which they began. [21] The poorly precisified similarity notion is replaced by the scientifically rigorous account of the mechanism *in virtue of which* members of a kind are similar. Indeed, an account of the mechanism is simply a (more) scientifically rigorous account of the relevant kind; it determines the comparative similarity with which kindhood is intertwined. It is in this way that Quine thinks it is "… a very special mark of the maturity of a branch of science

[20] As will be seen in section 4, this view fits nicely with Hacking's point that scientific inquiry allows our understanding of a kind to progress from a Russell-kind, to a Pierce-kind, and then to a Mill-kind.

[21] As will be seen in section 5, this eventual mechanistic account of a kind fits nicely with the underlying homeostatic mechanism Boyd requires to explain the clustering properties of a purported natural kind.

that it no longer needs an irreducible notion of similarity and kind."[22]

I agree with Quine on a number of key points. I mean to stay neutral, however, about his general story of learning, and think it likely that our folk conception of kinds at least sometimes results from processes *other* than induction. I do not, then, consider the generalizations offered by scientific inquiry as meta-inductions. Following Quine, however, I take kinds to be more or less natural depending on the success of science's generalizations. Most important, I embrace Quine's key insight that an account of a natural kind is most rigorously offered by an account of the mechanism responsible for the (comparative) similarity that delineates the kind's members. We need now to say more about these sorts of kinds and the mechanisms that inquiry into them uncovers.

4. Hacking's 4 Successive Conceptions

In an influential exchange between Ian Hacking and Richard Boyd concerning natural kinds, Hacking, liberally, identifies 4 conceptions of natural kinds: Russell-kinds, Mill-kinds, Pierce-kinds and Leibniz-kinds. While Hacking makes clear that he does not think we should be "bullied" into accepting only one conception, he does seem to think that each subsequent conception is rightly understood as enriching its predecessor.[23]

Hacking begins with Russell, understanding Russell-kinds to be those whose members have a number of properties "not known to be logically interconnected."[24] Future theorists gleaned much from these two features, but Russell-kinds are too weak. Kinds having only one

---

[22] Ibid., 174.

[23] In what follows, I will be less concerned with the accuracy of the respective authors' scholarship, and more concerned to make the notions of natural kinds they have tried to liberate from the tradition clear. All claims in this section about kinds according to Russell, for instance, should be understood as Russell-kinds according to Hacking.

[24] Hacking, "A Tradition of Natural Kinds," 112.

property in common (e.g., white things) or kinds who may not have any properties *all* in common (e.g., games) will be excluded, but a many kinds we would pre-theoretically reject as natural (e.g., white cats) will be included. The account could be rendered more exclusive by stipulating that the common properties are *natural* properties, but that is unacceptably question-begging in the current discourse. An effort to guarantee that natural kinds are suitably *natural* is, however, evidenced in the second distinguishing feature of Russell-kinds—that the common properties are not known to be logically interconnected. The purpose of this criterion is to restrict natural kinds to those that share properties that we observe for reasons that we do not. Though Hacking spends limited time on it, this criterion is intended to block categorizing kinds *merely* of our own devising as natural. An attempt to block the inclusion of kinds *merely* of our own devising is offered by all natural kind theorists.

Mill-kinds include supplements to both of Russell's key ideas. Natural kinds are held to be "real Kinds," which are appropriately contrasted with "finite kinds." The real Kinds are those with not just *a* number of properties, as Russell had it, but a "plausibly inexhaustible" number. It is not clear how many common properties are required for Hacking's Mill to consider their inexhaustibility plausible. What is clear is that real Kinds, i.e., natural kinds, cannot simply share set of *specified* common properties thereby *known* to be finite. In this way, Mill also seems to incorporate Russell's second key idea.

Hacking's explication of Mill emphasizes an additional requirement, not explicit in Russell, that members of a natural kind constitute the maximal class having the plausibly inexhaustible number of properties. The virtue of this supplement is that it excludes the troubling cases still included by Russell's stated criteria—e.g., white cats. Mill-kinds include only those of suitable generality. For this reason, the maximal class condition seems to have been adopted by

every natural kind theorist following Mill.

While Mill-kinds supplement both of Russell's key ideas and add a third, Pierce-kinds reverse Russell's requirement that the logical connection responsible for the shared properties be unknown. As Hacking's Pierce would have it, one of the tasks of science is to identify the laws *in virtue of which* members of a kind share their properties *and in which* reference to the kinds is featured. We identify cats as cats in virtue of a number of shared observable properties, and then expect the scientist to identify *why* cats share the properties they do and *what laws* are true of cats. On Hacking's construal then, a Pierce-kind is a class of objects with a potentially inexhaustible number of properties in common "… but such that there is a systematized body of law about things of this kind, and is such that we may reasonably think that it provides explanation sketches of why things of this kind have many of their properties."[25] Interestingly, Pierce's motivation seems to be the same as Russell's: exclude kinds merely of our own devising. Until we have discovered the unobservable laws governing the kind and how they explain the observable properties, we cannot legitimately pronounce the kind natural—we may, that is, *simply* have made it up. Such kinds, as Russell and Mill would agree, would not be natural.

Not only can Pierce-kinds and Mill-kinds be thought of as supplementing Russell-kinds, but scientific progress in *understanding* a kind is sometimes a move from an understanding of the kind as Russellian, to Millian, and finally to Piercian. This last step is emphasized by Hacking, who takes it as a truism that natural science (at least sometimes) progresses along these lines. Quine's contention that progress in understanding a kind is made by refining the relevant comparative similarity points to the same phenomenon. We start with a notion of a kind with a number of shared properties in common. Upon investigation, we may discover that the number

[25] Ibid., 120.

of shared properties is potentially infinite, though remaining ignorant about why. Upon yet further investigation, we may discover laws true of those kinds and provide explanations of why members of the kind share the manifold properties that they do. This progress provides new, and more successful, generalizations about a kind's members. When this happens, the purported kind will have progressed from a Russell, to a Mill, to a Pierce-kind. Just as Quine warned that we should not *require* that natural kinds originate as non-theoretical, however, so Hacking points out that sometimes we develop a notion of a Pierce-kind as a result of accepting natural laws entailing their existence.

One may reject that science is in the law business altogether, and Leibniz-kinds cover that contingency. Leibniz-kinds are identified as featuring not in laws, but in models. Here it is not the requirements for the *naturalness* of kinds that are supplemented *per se*, but the view of the generalizations that feature references to natural kinds. According to Hacking's Leibniz, science builds models that specify the underlying structure of natural kinds. Kinds are not, therefore, appropriately considered natural when they are referred to in relevant laws, but when a posited underlying structure of the kinds supports the construction of a predictive and explanatory model. The potentially inexhaustible number of properties shared by members of the kind is not understood as a function of their governance by laws true of those members, but of what can be *inferred* from the members' common underlying structure. In Hacking's words, Leibniz-kinds are "… classes for which there are associated known structures from which can be inferred properties in common among objects of that class."[26] A kind is natural when an identified, underlying structure *explains* the observable manifold properties.

If the account I have been developing is correct, then the requirement that to be natural a

[26] Ibid., 121.

kind *must* feature in exceptionless laws of nature is gratuitous. What science does is an important and interesting philosophical question. But *whatever* science does—whether it constructs models as Leibniz would have it, offers meta-inductions as Quine would have it, or discovers laws as Pierce would have it[27]—which kinds are natural is an independent question. Science is doing *something* that delivers generalizations—regardless of the status of those generalizations. Call the work of science 'Sam'. If a kind and the term used to refer to it is useful for Sam, then refusing to bestow the honorary 'natural' on that kind and its term is illegitimately prejudicial.

From Hacking, then, we are given four conceptions of natural kinds: Russell-kinds, Mill-kinds, Pierce-kinds, and Leibniz-kinds. Indeed, he thinks that we often use the same term—'natural kind'— to denote one or the other of these kinds of kinds in "different discourses." These notions of natural kinds are similar enough such that "[u]sually it makes not the slightest difference [which kind we mean], and if there is a possibility of confusing the slightly different extensions, we make plain what we are talking about."[28] In terms of utility, the point is that sometimes its useful to talk about Russell-kinds, sometimes Mill-kinds, and sometimes Pierce or Leibniz-kinds—our opinion on this latter utility depending on our preferred view of the deliverances of science. No one of these notions is, on Hacking's view, to be privileged above the other, though each can be seen as supplementing the former. Despite Hacking and Boyd's disagreements about the history of the natural kinds tradition, Boyd's homeostatic cluster theory of natural kinds is welcomed by Hacking as yet another supplement.

5. Boyd's Homeostatic Clusters

Boyd influentially identifies natural kinds as those whose members have clusters of properties

[27] Or whether these views are actually compatible, as I suspect.
[28] Ibid., 120.

that usually co-occur as determined ("brought about") by an underlying homeostatic mechanism. Boyd takes the role of kinds in explanation and induction to be their central function and recommends looking to actual scientific inquiry to determine which kinds are useful for those purposes.

Three conditions, according to Boyd, will be true of any such kinds:

> First, the definition of the kind in question is determined *a posteriori* rather than by social convention. Second, this in turn is true because the use we make of reference to the kind in induction and explanation requires that it be defined by a set or cluster of properties whose membership is determined by the causal structure of the world and is thus, in a relevant sense (which I propose to specify) independent of our conventions or our theorizing. Finally, the establishment of a relation of reference between a term and the kind in question depends on the existence of an epistemically relevant pattern of causal relations between instantiations of the kind and the use of the term in question *and not on the acceptance by any or all of the members of the relevant linguistic community of the correct definition of the kind*. ("Realism, Anti-Foundationalism and the Enthusiasm for Natural Kinds," 129.)

The first condition is the most straightforward, and the second provides its defense. Natural kinds are not determined *a priori,* or, more germane to Boyd's concerns, as a consequence of any foundationalist principles. The second condition requires natural kinds to be defined by a usually co-occurring cluster of properties, whose clustering is "determined by" the causal structure of the world. The opaque "relevant sense" in which this determination is independent of our theorizing is that the mechanism in virtue of which the members share their manifold properties be *causal*—where causal is contrasted with conceptual. It is not that we do not conceptualize. Rather, it is that the clustering of the manifold properties is not *in virtue* of those conceptualizations, but in virtue of the world's causal structure. Indeed, Boyd (1991, 1999) requires further that the causal mechanism be homeostatic; external or internal pressures that keep the properties clustering together must be operative on the mechanism. If there is no *causal* mechanism determining the property clustering, then the *clustering* of the properties is held to be

merely a result of our conceptualization. Such kinds are not suitably natural—it is unclear whether we would even get *kinds* as against mere sets. Since the mechanism in virtue of which the properties cluster is causal and not *merely* conceptual, kind membership must be determined *a posteriori*, and it is in this way that Boyd's second condition supports the first.

That the mechanism responsible for the clustering of properties be a *homeostatic* mechanism deserves another word here. As above, for a mechanism to be homeostatic is for it to be subject to pressures that keep the various properties it determines clustering together. Boyd's paradigm of a natural kind is species and is argued to exhibit homeostasis; there are internal pressures (genes) and external pressures (environment) that keep the clustering properties of any given species clustering. The nature of the requisite pressures, however, is unclear. Taking the lead from Boyd, we may take the properties to be *causally* related.[29] Is it required, then, that the properties in the cluster be under pressure such that they are causally *dependent* on one another or will mere causal *relevance* do?

Though I think this question deserves to be taken up at length in its own right, its answer is irrelevant for my arguments for the non-naturalness of pain. If it is required that the properties are causally *dependent* on one another, then the ubiquity of the dissociations between the properties in the paradigmatic property cluster for pain will show that there this requirement is unmet. As we will see most clearly in chapter 2, too many of the properties are present in the absence of the others too often for the claim that they are causally *dependent* on one another to be credible. If the requirement is weakened to mere causal relevance, however, then much more than the properties clustering in paradigmatic *pain* experiences will be included. As we will see

[29] Mechanisms that ensure property clustering in some non-causal fashion would, I take it, simply not be stringent enough to yield a kind that is the result of the causal structure of the world, as against merely the structure of our concepts, as Boyd requires.

most clearly in chapter 3, the underlying mechanism that would determine all the properties causally *relevant* to those involved in paradigmatic pain experiences would be the brain in its entirety, and perhaps even the entire central nervous system.[30] If we require causal dependence then some properties will get left out of the cluster, but if we require causal relevance then too many mechanisms will be included.

Boyd's third condition is that a term's reference to a natural kind depends on establishing "epistemically relevant patterns of causal relations" between the kinds' members and the natural kind terms and *not* on acceptance of the definition of the terms by the relevant linguistic community. The epistemically relevant patterns of causal relations here are that the term facilitates either explanations or inductions about members of that kind. This is consistent with Boyd's insistence that the core of a natural kind is its successful role in explanation and induction. This third condition, then, is best understood as making clear that the relevant linguistic communities definition of the kind is decidedly *not* a condition.

English speakers, for example, may not accept that cloves of garlic are lilies, that tomatoes are fruit, or that whales are fish. The linguistic community is more or less likely to use 'fruit' to refer also to tomatoes, 'fish' to refer also to whales, and 'lilies' to refer also to garlic. Whether these natural kind terms *do* so refer when employed in the generalizations of the relevant scientific discourses, however, is not beholden to that acceptance. Instead, whether the natural kind of 'lilaceae' includes garlic turns on establishing that, and only that, using the term 'lilaceae' to refer also to garlic facilitates better inductions and explanations pertaining to lilies and garlic than otherwise. The acceptance of the linguistic community is irrelevant.

---

[30] I am indebted to Michael Devitt for pointing out that without *some* further tightening of "causal relevance" much more than just the nervous system will be included, as there seems little limit to what may be construed as causally relevant to anything else.

Continuing on in the tradition identified by Hacking, Boyd-kinds supplement Leibniz-kinds in three major ways. Recall that Leibniz-kinds are those kinds forming a maximal class with a plausibly inexhaustible set of properties inferred from some identified structure that supports the construction of a predictive and explanatory model. The first supplement from Boyd is his further specification of the *identified structure* as an underlying homeostatic mechanism. Second, Boyd *weakens* the plausibly inexhaustible number of properties to a clustering of properties that usually, but only usually, co-occur. Finally, Leibniz's posited explanatory relations between the model and the world in general are *tightened* by Boyd as holding specifically between the use of terms in generalizations and the kinds those terms refer to when employed in the generalization.

These supplements have made Boyd-kinds the richest account of natural kinds to date. It would nonetheless be unwise to privilege Boyd-kinds as the *only* legitimate conception of natural kinds. As Hacking points out, there is just no good reason to privilege one conception to the exclusion of all others. Note that this is because each of the different conceptions yields a notion of kind that may fulfill the scientific utility criterion for different inquiries at different times. So too, as a science progresses, we ought to be able to develop richer and richer conceptions of the kind referred to in its generalizations. Just as Quine pointed out. If a kind is natural, it ought to eventually be understood to be, *or at least not be shown unable to eventually be understood to be*, a Boyd-kind.

Rather than take Boyd's conditions as necessary and sufficient conditions, then, I will instead liberate 4 main indicators of the naturalness of a kind from Boyd's rich conception:

1. Natural kinds have, and are identified by, clusters of usually co-occurring properties.
2. Natural kinds have an underlying, homeostatic mechanism (that determines the cluster of

usually co-occurring properties).

3. Natural kinds feature in science's generalizations, such that the referring terms in those generalizations refer to kinds that the correlative generalizations featuring those terms explain.
4. Natural kinds are usefully referred to for explanation and induction (because of indicators 1-3); i.e., insofar as reference to a kind is *not* useful for induction and explanation, the kind is not natural.

Kinds, I contend, are more or less natural insofar as they have these features.

In the chapters that follow, I will argue that the central reason pain is not a natural kind is because the usually co-occurring properties appropriately used to identify pain are not determined by an underlying homeostatic mechanism (indicator 2), and that reference to pain is not useful for induction and explanation (indicator 4). There is a cluster of usually co-occurring properties (indicator 1) by which we identify pain, but pain's failure to adequately exhibit indicators 2 and 4 reveals that the clustering of properties is one merely of our own devising. Pain is non-natural.

The two shortcomings of pain's naturalness are related. There is no *mechanism* (homeostatic or otherwise) underlying the properties usually co-occurring in paradigmatic pain experiences. Recall that, for a natural kind, it is the underlying mechanism that ensures both that the inductions and explanations concerning the clustered properties are adequately successful and that the properties exhibited by members of the kind are not simply sets constructed by a linguistic communities' conceptualization of the member's observed common properties. When there is no underlying mechanism that explains the properties we use to pick out the phenomenon, as I will argue is true in the case of pain, all we have is the conceptualization of the

folk. This is insufficient to render the kind natural—and results in unacceptably unsuccessful inductions and explanations. The reason pain is not useful when referred to in science's generalizations— as the scientific utility criterion requires of natural kinds— is that there is no pain mechanism.

One might worry that if this much is required for a kind to be natural we need not be much concerned if a kind is non-natural.[31] Functionalism, in various guises, is widely accepted as the best approach to the mind. Many functionalists, moreover, hold that functionally-characterized kinds are multiply realized. Multiple realization may be thought to undermine the need to identify a *single,* underlying homeostatic mechanism; a functionally-characterized mental phenomenon may be explained by a *multiplicity* of mechanisms. If a single mechanistic explanation is required for natural kinds, then the whole mind is non-natural—which is surely an unhappy result. Chapter 5 addresses this and related concerns in detail when discussing pluralism, but a brief word is warranted now.

One needs to distinguish the claim that there is *no* mechanistic explanation of a phenomenon for a type of creature from the claim that there are *many* mechanistic explanations of a phenomenon for different types of creatures. That a kind is realized by different mechanisms when realized by different organisms need not undermine the naturalness of a kind. *That a kind does not submit to a mechanistic explanation when realized by a single type of organism, however, is another matter.* As Putnam himself, a champion of multiple realizability, points out:

> …to investigate this [functional] hypothesis is just to attempt to produce 'mechanical' models of organisms—and is not this, in a sense, just what psychology is about? The difficult step, of course, will be to pass from models of *specific* organisms to a *normal form* for the psychological description of organisms. ("The Nature of Mental States," 200.)

[31] I am thankful to Michael Devitt for pressing this worry.

The claim that mental kinds were functionally characterizable and multiply realized was not intended to rule out the mechanistic explanations offered by scientific inquiry. Quite the contrary. Putnam, at least, was clear that the functional characterization of a mental kind was intended as an important step in providing mechanistic explanations of it—though indeed allowing for different mechanistic explanations of a single phenomenon differently realized in different types of organisms.

Similar considerations apply to the rather different type of functionalism championed by David Lewis. While the causal roles enshrined in the folk platitudes about a type of mental state may have various occupants in various kinds of creatures, there is presumed to be some one state that plays that causal role for (at least most of) the creatures in a population. Lewis requires that population themselves are natural kinds. Thus of pain in particular, Lewis says:

> Human pain is the state that occupies the role of pain in humans… A state occupies a causal role for a population and the concept of occupant of that role applies to it, if and only if, with few exceptions, whenever a member of that population is in that state his being in that state has the sort of causes and effects given by the role. ("Psychophysical and Theoretical Identification," 231.)

Putting aside the well-known differences between Putnam and Lewis's functional characterizations of mental kinds, both types of functionalism leave open, indeed assume, that the "realizers" or "occupants" of functionally-characterized kinds will be identified for a "type of organism" or "for a population." My claim is that if those mental kinds are usefully referred to in science's generalizations, the relevant identifications will be of a mechanism that explains the clustering properties of the kind. Noting, of course, that if one is a believer in multiple-realizability, many such identifications are expected, though only one for, say, humans. Functionalism is not in any conflict with the claim that we need a mechanistic explanation of a kind for reference to it to be useful for scientific inquiry—that is, for the kind to be natural.

In the following chapters I argue that there is *no* mechanistic explanation for pain or any pain type. There is no pain mechanism for humans—and I see no reason to think there is one for tigers, Martians, or computers, but that is irrelevant here. Though there are many mechanisms implicated in pain experiences, no one of them determines the pain experience of any creature, and as such, neither pain nor any pain "type" is a natural kind.

To evaluate this claim, more needs to be said about mechanistic explanations and how they relate mechanisms to the phenomena they bring about. Section 6 provides a brief overview of mechanisms and mechanistic explanation, and their relationship to Boyd-kinds.

6. Of Mechanisms

Mechanisms and mechanistic explanation are currently receiving renewed attention in the philosophy of mind and philosophy of science. In this section, I provide a brief sketch of insights from three major theorists in the contemporary mechanism literature: Glennan, Bechtel, and Craver.

While all three theorists have many disagreements, their understandings of mechanisms share much in common. On Glennan's view, a mechanism is always defined *for* a behavior, and "[a] mechanism for a behavior is a complex system that produces behavior by the interaction of a number of parts, where the interactions between parts can be characterized by direct, invariant, change-relating generalizations."[32] According to Bechtel, similarly: "A mechanism is a structure performing a function in virtue of its component parts, component operations and their organization. The orchestrated functioning of the mechanism is responsible for one or more

[32] Glennan, "Modeling Mechanisms," 445, and Glennan, "Rethinking mechanistic explanation," S344.

phenomena."[33] Finally, according to Craver: "Mechanisms are entities and activities organized such that they are productive of regular changes from start or set-up to finish or termination conditions."[34]

For purposes here, it is useful to emphasize the commonality among these accounts: Mechanisms are wholes with *parts* that are *organized* and *do things*, such that the organization and goings-on (interaction, operations, activities) of the parts (entities) determines (produces, is responsible for, are productive of) some behavior or phenomenon of interest. According to these theorists, while there may also be other kinds of explanations, if the mechanism determining a phenomenon of interest is appropriately identified, then that phenomenon has been explained.[35] I consider any representation of a mechanism, however partial, to be a mechanistic model,[36] and any provision of a mechanism that purports to determine some phenomenon of interest a mechanistic explanation.

Note that these mechanism theorists claim that mechanisms determine *phenomena* or *behavior*, whereas Boyd-kinds require underlying, homeostatic mechanisms that determine *a cluster of properties*. These two accounts are easily seen to be harmonious. Boyd is focused on natural kinds and, consistent with that tradition, looks to identify underlying homeostatic mechanisms bringing about the clustering properties we use to identify a kind. When the

---

[33] Bechtel and Abrahamsen, "Explanation: a mechanist alternative," 423. For the initial offering and elaborate defense, see Bechtel and Richardson (1993). For application and further explication, see Bechtel (2008).

[34] Machamer, Darden, and Craver, "Thinking about mechanisms," 3. For further elaboration, see Craver (2006). For application and further explication, see Craver (2007).

[35] There are interesting questions concerning whether the *mechanism* explains, as Craver (2007, 27) and Glennan (2008, 448) would have it, or whether it is the representations (models, as I will call them) of mechanisms that are explanatory, as Bechtel would have it (2008, 17-19). Nothing should turn on this disagreement for present purposes, but I think the good money is on Bechtel.

[36] I gloss over Craver's distinction between explanatory sketches and schemas and the corresponding how-possibly, how-plausibly, and how-actually models.

phenomenon we are interested in is a mental phenomenon, as in the case of pain, identifying the kind by the clustering properties will be equivalent to identifying the kind by the relevant behavior or observable phenomena. We identify the mental kind of interest by some truckload of properties that characterizes it.[37] For Boyd, one indicator that the identified kind is a *natural* kind is that the properties by which we identify the kind are determined by an underlying, homeostatic mechanism. For Glennan, Bechtel, and Craver, we will have explained the mental, natural kind when we have identified the mechanism that determines the phenomenon of interest. Recall further that determination of the underlying mechanism is, according to Boyd, what allows for fruitful explanation and induction, since it is the underlying mechanism that is firmly situated in the world's causal structure. Glennan, Bechtel and Craver are focused on how mechanistic models[38] explain, and one virtue of a good mechanistic explanation is that it yields successful predictions. The similarity here should be obvious. The mechanisms of interest to the mechanism theorists are those that explain some phenomenon of interest and yield predictions, and the mechanism of interest to Boyd are those that provide for fruitful inductions and explanations.

Though one needn't accept Boyd's account of natural kinds to accept any of the mechanism theories, if we *do* accept Boyd-kinds as natural kinds, then we may understand Glennan, Bechtel, and Craver as supplementing Boyd's account by filling in *how, why, and when* underlying homeostatic mechanisms yield fruitful inductions and explanations. The recent work by these theorists provides a mechanistic update to Boyd's theory of natural kinds.

While opinions among mechanism theorists vary concerning when a mechanistic

---

[37] This is not to say that the mental kind just *is* equivalent to behavior or observable phenomena—only that it is by some observable properties or other that we *identify* the mental kind for further inquiry. This is so even if, as some hold (wrongly, I think), that mental states are utterly private. *Identification* of mental state types would still involve observation by the person in the occurrent state that is a member of the identified kind.

[38] Or, again, according to Craver and Glennan, the mechanisms themselves.

explanation is a *good* one, all agree that one key requirement is an appropriate characterization of the phenomenon to be explained. Craver[39] usefully identifies 3 ways this may go wrong. First, one may inquire about the mechanism of a fictional phenomenon.[40] If there's no such thing as phlogiston, for example, then an appropriate characterization of it will prove intractable and, accordingly, so will providing a good mechanistic explanation. Second, one may commit one of two major taxonomic errors, which Craver refers to as *lumping* errors and *splitting* errors.[41] As their titles suggest, a lumping error is when several phenomena are incorrectly characterized as one (e.g., memory, he thinks) and a splitting error is when one phenomenon is incorrectly characterized as many (e.g., oxidation antecedently understood variously as rusting, burning, and breathing). Lastly, one may under-specify the phenomenon.[42] If one under-specifies the phenomenon, then there is a danger that the mechanism one posits to explain it will also be underspecified. An underspecified mechanism would be unable to account for every aspect of the phenomenon to be explained. As Craver puts the point:

> One can conjecture a mechanism that adequately accounts for some narrow range of features of the phenomenon, but that cannot accommodate the rest. For this reason, descriptions of the multiple features of a phenomenon, of its precipitating, inhibiting, modulating, and nonstandard conditions, and of its by-products, all constrain mechanical explanations….[43,44]

---

[39] C.f. Glennan (2005) and Bechtel (2008). I have chosen to focus on Craver's view of characterizing the phenomenon here solely because of his clear identification of 3 major ways it can go wrong. Though there is perhaps some difference in emphasis, all three theorists require adequate identification of the phenomenon to be explained, and I see no reason to think that the pitfalls identified by Craver would not be agreed to *be* pitfalls by both Glennan and Bechtel.
[40] Craver, *Explaining the Brain, Mechanisms and the Mosaic Unity of Science*, 122.
[41] Ibid., 123-124.
[42] Ibid., 127.
[43] Ibid., 128.
[44] Note Bechtel's similar point (2008, p. 61): "Failing to differentiate phenomena that are really different from each other can result in proposed explanations that fail to adequately account for the full character of the phenomena."

If one under-specifies the phenomena, then the corresponding mechanical model will, correlatively, "under explain" the phenomenon.

In the following three chapters I argue that in the case of pain we have a dilemma: The specifications of the phenomenon inevitably either under-specify or lump. Paradigmatic pain experiences, I will argue, include several aspects, and exclusively privileging any one of these aspects yields an inappropriate characterization of the phenomenon. Views that isolate some one aspect commit the error of under-specification. If, on the other hand, one appropriately identifies the phenomenon of pain using the full cluster of its usually co-occurring properties, then, I will argue, there is no one underlying, homeostatic mechanism that determines them. The different aspects of pain are not determined by one underlying mechanism. The folk notion of pain lumps many phenomenon into one—it is a kind merely of our own devising. There is no appropriate specification of pain that will yield a single phenomenon determined by an underlying, homeostatic mechanism—all accounts commit either under-specification or lumping errors. These problems are evidenced in the recent move in pain science towards providing a mechanism-based characterization of pain.

7. A mechanism-based classification of pain

Until recently, pain research and diagnostics exclusively employed a model of pain focused on pathologies. A patient presenting with pain is, according to this model, to be diagnosed such that the presumed pathological origin of their pain is identified and, if possible, eliminated. This approach is called the disease, or medical, model. While the medical model may seem obvious or even inevitable, it has recently come under scrutiny by a number of prominent pain scientists.

These skeptics have begun turning, instead, towards a mechanism-based classification of pain.[45] This alternative approach seeks to classify pains based on whatever underlying mechanism determines them, and has been dubbed 'pain analysis.'

In a landmark editorial in *Pain*, Woolf, et al (1998) present the idea that the time in pain research had come when the medical/disease model should be replaced with a mechanism-based classification of pain that now could, and should, be sought. Smart, et al (2008) helpfully elaborates the traditional medical/disease approach and its problematic assumptions:

> The medical/disease model refers to a tissue and pathology oriented approach towards the explanation of pain and dysfunction and functions on the premise that all pain has a dominant tissue or structural source. … The medical model is based on three broad assumptions:
> - There is a direct and proportional relationship between the amount of tissue damage and the level of pain experienced.
> - There is a linear relationship between the amount of tissue injury and the degree of any ensuing physical impairment or disability.
> - Treating the pathology resolves the pain and therefore the disability.
>
> … [I]nterpretations and classifications of pain based on the medical/disease model and its structure-oriented approach have been criticized since such approaches do not explain the complexity and variability of pain presentations observed clinically; where pain is observed to persist after healing, where patients' reports of pain in response to similar injuries differ greatly or where the magnitude of a pain report does not match the injury or pathology. Nor do they account for the variations in patients' outcomes in response to interventions directed towards its treatment since by the tenants of the traditional medical model treatment of the pathological process should be accompanied by relief of pain (Smart, et al, 2008).

The motivation for the move to a mechanism-based classification of pain is, I hope, clear. There are patients with the same pathology, but varying pain, and (at least) similar pains but varying pathologies, and pains with seemingly *no* pathology. We need, of course, to eliminate the pathologies, but we also need to eliminate the pain. As Woolf's initial editorial forcefully argued, we need a characterization of pain that will explain it more generally than the disease model

[45] Jensen and Baron, 2003, Finnerup and Jensen 2006, Smart et al 2008, Woolf et al, 1998, Woolf 2004, and Voisin 2001, are notable, but by no means exhaustive, examples.

allows. He urged the scientific community to look for general *mechanisms* responsible for pains whenever we find them—even, and especially, when they are not correlated with a disease, pathology, or tissue damage. The medical/disease model neither offers, nor even allows, for identifying *pain* mechanism(s). The cry of the movement was that we need a new model.

From the beginning of the movement towards mechanism-based classifications of pain, including the influential 1998 editorial, different *types* of pain explained by what were already then recognized as *different* mechanisms were acknowledged. The reason is straightforward. Though there is disagreement about how pains should be grouped, if the groupings are to correspond to a single, underlying mechanism, then different types of pain are going to have to be grouped separately. Neuropathic pain and tissue injury pain, for an obvious instance, are known *not* to be determined or maintained by the same underlying mechanisms. Woolf, et al (1998), initially proposed 3 broad types of pain—transient pain, tissue injury pain, and nervous system injury pain—and corresponding possible mechanisms. For the latter two types, multiple possible mechanisms were identified (7 for tissue injury pain and 8 for nervous system injury pain). Other, similarly rudimentary, mechanism-based classifications of pains have followed.[46] The new mechanism-based models do not seek to identify *the* pain mechanism; rather, they seek to identify the various mechanisms explaining various types of pain. A mechanism-based classification of pain, while *more* general then disease or pathology-based classifications, will still require divisions among types of pain. For this reason, some pain scientists supporting a mechanism-based approach to pain suggest that mechanism-based classifications should be understood as a *supplement* to the traditional medical/disease model.[47] Some token pains *are* explained by an underlying disease and pathology, such that elimination of that pathology

[46] For instance: Arner (1998), Gifford (1997), Jones (1995), and Lidbeck (2002).
[47] The possibility is in fact raised by Woolf, et al (1998), in the initial paper.

eliminates the pain; perhaps "tissue-damage pain" should simply count as one among many types of pain. Mechanism-based classification approaches were intended, from their inception, to eventually classify *different types of pain* corresponding to *different* underlying mechanisms.

To identify the underlying mechanism responsible for any pain type, scientists and physicians must use the symptoms and signs presented by patients. In particular, they must use patients' reports. As Finnerup and Jensen (2006), point out: "Because patients report symptoms rather than mechanisms of pain, the main clinical problem is to extrapolate mechanisms from symptoms or clusters of symptoms."[48] The failed attempts to provide translations of symptoms and signs into mechanisms evidence the enormity of this problem.[49] Even limiting the classification schema to distinct types of pain leaves double dissociations between symptoms and any mechanism proffered as candidates for explaining them. These problems are recognized by the mechanism-based advocates—and are recognized as being as yet intractable.

These points are summarized nicely in Finnerup and Jensen's (2006) critical analysis of extant mechanism-based classifications of neuropathic pain:

- To develop an effective mechanism-based classification, we need to be able to relate symptoms and signs to mechanisms and to identify specific treatments for specific mechanisms.
- One mechanism might give rise to different symptoms or signs; for example, upregulation of sodium channels in C-fibres increases fibre activity, resulting in burning pain, paroxysms and dynamic mechanical allodynia.
- Similarly, one symptom or sign can be caused by several initiating mechanisms; for example, cold allodynia can be attributed to different mechanisms in peripheral and central neuropathic pain.
- It is becoming clear that we have not yet obtained a viable mechanism-based classification for neuropathic pain, and more-rigorous studies are required to test this approach (113).

---

[48] Finnerup and Jensen, "Mechanisms of Disease: mechanism-based classification of neuropathic pain—a critical analysis," 110.

[49] See, for instance, Jensen and Baron's important "Translation of symptoms and signs into mechanisms in neuropathic pain," 2003, for a rich attempt.

Note that this is a summary of the key points in a critical review of mechanism-based classifications of a particular *type* of pain— neuropathic pain. *For that purported type of pain* there are double dissociations between known mechanisms associated with neuropathic pain and neuropathic pain symptoms. The double dissociations between symptoms and signs and underlying mechanisms, even *within* particular pain types, are the largest problem facing the new movement towards a mechanism-based classification of pain.

The claim that pain is not a natural kind can now usefully be divided into two distinct claims. *First, the many different types of pain do not constitute a natural kind.* There is no underlying homeostatic mechanism that explains all pain. If we are going to have a mechanistic explanation for pain, then we are going to have to group the pains explained by an underlying mechanism into several types. *Second, however, no type of pain is itself a natural kind.* Even once we separate distinct types of pain, there is no underlying mechanism that explains any one of them.

Using the account of natural kinds developed in the previous sections, we can understand the mechanism-based approach to pain classification as an attempt to render pain a natural kind. The approach is intended to explain clusters of properties (signs and symptoms) by their determination by underlying (presumably, homeostatic) mechanisms, with the explicit hope that doing so will be of greater utility in diagnosis and treatment. Note, however, that even the sought-for success would only render particular *types* of pain natural kinds. My first claim then, that the different types of pain do not constitute a natural kind, is one that the pain scientists taking a mechanism-based approach seem themselves to be committed to—though perhaps without realizing it as such. My second claim, however, that no type of pain is itself a natural kind, is one that even those pain scientists taking a mechanism-based approach are not (yet)

willing to grant. They recognize that as of now no particular type of pain can be explained by an underlying mechanism, but remain hopeful that such explanations are in the offing.

I have said that I argue in the following chapters that the attempts at a mechanistic explanation for pain either under-specify the phenomenon or commit the taxonomic error of lumping, but before turning to these arguments, note these two errors in the motivation and problems for the mechanism-based movement in pain science. The approach is motivated by the perceived failure of the disease/medical model to explain pain. Mechanism-based advocates assert that the model should be rejected because of its overly specific identification of pain as determined by particular diseases, pathologies, and tissue damage. That is, the traditional model *under-specifies* pain. On the other hand, the biggest problem facing mechanism-based classifications of pain is the double-dissociations between symptoms and mechanisms, even for particular pain *types*. This problem, I will argue, is insurmountable. It is the result of the lumping error that treats several aspects of paradigmatic pain experiences as a phenomenon determined by a single, underlying homeostatic mechanism. The mechanism-based approach is failing because it lumps and the traditional model has failed because it over-specifies.

Though pain is real and important for everyday life, neither pain nor any type of pain is a natural kind. Natural kinds are those that are usefully referred to in the generalizations of the relevant science(s). Boyd's offered criteria for Boyd-kinds provide useful indicators for when a kind is natural, and pain fails to adequately meet those criteria. This is so because there is no single, underlying homeostatic mechanism that determines the clustering properties of either pain or any type of pain. It is to the arguments for these conclusions that I now turn.

Chapter 2: The Inadequacy of Unitary Accounts

1. Introduction

Arriving at an adequate account of the phenomenon is the first step in determining whether pain is a natural kind. Though pain scientists now generally understand pain to be a complex experience involving sensation, emotion, cognition, *and* motivational responses, most philosophers maintain that pain is best characterized by *one* aspect of this complexity. I call all such accounts unitary accounts. In the following two chapters, I argue that while complex accounts do better justice to the phenomenon of pain, the resultant kind is non-natural. In this chapter, I argue that no unitary account of pain can accurately characterize the phenomenon.

There are three basic types of mental states often focused on in philosophical, unitary, accounts of pain: perceptions, emotions, and sensations. There is, however, much disagreement about the phenomena each of these terms is best used to identify. My goal is to put questions of usage aside and focus on which, if any, kind of mental state pain might be appropriately characterized as—whatever we call it. To this end, I lay out rough definitions of these types of mental states intended to be as broad and uncontroversial as possible.

Sensations, perceptions and emotions are here characterized by the way in which they vary across two kinds of content: intentional content and qualitative content. Emotions and perceptions have intentional and qualitative contents, while sensations have only qualitative content.[50] Emotions are here characterized as having characteristic qualitative and intentional contents, which differ in kind from the contents characteristic of perception—though I remain

[50] While I favor talk of qualitative *character* as against qualitative content, I will use the term 'content' for a more neutral topography.

neutral about just what the characteristic contents of emotions are.[51] I will understand intentional content to be the *aboutness* of the mental state; a mental state has intentional content when that state is *about* something. Qualitative content will be here understood to be the *feeling* of a mental state; a state's qualitative content will be characterized by what it feels like to have it. While both feelings and qualitative content are held by some to be always conscious, I will make no such assumption here. Similarly, the terms 'quale' and its plural 'qualia' are often used to refer to qualitative content(s) by those who hold that they, or the states that have them, are always conscious. My use of those terms is not intended to signal any such commitment. Whether qualia or states with qualitative content are always conscious is a vexed issue irrelevant for purposes here.

Philosophers have identified pain with each of the three kinds of mental states characterized above: sensations, perceptions, and emotions.[52] In this chapter, I examine each of these identifications in turn. Section 2 focuses on the theory that pain is a sensation, as presented most influentially by Kripke (1980). In Section 3 I examine the idea that pain is a perception, as presented most influentially by Armstrong (1962) and Tye (2000, 2005). In section 4, I briefly discuss the relatively new idea that pain is an emotion, as presented by Clark (2005), Gustafson (2005), and Craig (2003a, 2003b, 2003c). Consideration of the emotion views reveals that these theorists actually take pain to be a complex experience of which emotion, as I have characterized it, is only a proper part. The richest incarnation of this view is presented in Craig and will be

---

[51] Note that if emotions are characterized perceptions of bodily states, then emotions are simply a kind of perception, though nonetheless with characteristic content.

[52] I do not discuss the more recent characterizations of pain as an imperative. As these accounts are neither traditional, nor (yet) very influential, nor consistently unitary, I put them aside for purposes here. For a clear presentation of the imperative view of pain see Klein (2007), and for a rich discussion of the problems with either a complex *or* unitary account of pain as an imperative see Bain (2011).

taken up in more detail in Chapter 3 with the other complex accounts. I conclude in section 5 with some general remarks about why no unitary conception of pain is likely to succeed.

Each of the unitary views suffers from identifying some one aspect of pain as necessary and sufficient, when consideration of actual cases of pain reveals that no one aspect is both or either. It is worth noting at the outset that I do *not* deny that paradigmatic pain experiences have sensation, emotion, and perception as components. Indeed, on my view, paradigmatic pain experiences also have thoughts and motivational responses as components. A typical pain feels like something, is about something, includes a perception of something, and makes us want to do something. The thesis of this chapter is only that pain is not properly identified with any *one* of these: an adequate characterization of pain acknowledges all of its paradigmatic components.

2. Pain is Not Simply a Sensation

As sensations are those mental states having only qualitative content, the claim that pain is simply a sensation is equivalent to the claim that pain is a type of mental state exhausted by its qualitative content. Saul Kripke's 1972 *Naming and Necessity* lectures provide one of the most famous presentations of this view. Kripke's position is that there is a distinctive qualitative content essential to and exhaustive of pain. I argue in this section that for any identified qualitative content, there are token pains without it; though pains usually have qualitative content, there is no distinctive, qualitative content essential to pain. Furthermore, as pains may vary in qualitative content, if there is anything that constitutes all the tokens as a type, it must be

something *other* than their qualitative content. There is no distinctive, qualitative content exhaustive of pain.[53] Pains are not simply sensations.

Pain arises, in Kripke's lectures, as an application of his now well-known argument against identity theories. We are held to know *a priori* that true identity statements are necessarily true. Mind-brain correlations, however, are never necessary identities since we can always, Kripke claims, imagine any mental state occurring without its corresponding brain state. Mind-brain correlations are then merely contingent correspondences at best, and the purported identities are all false. In a passage quoted often in the pain literature, Kripke (1980) argues:

> In the case of molecular motion and heat there is something, namely, the sensation of heat, which is an intermediary between the external phenomenon and the observer. In the mental-physical case no such intermediary is possible, since here the physical phenomenon is supposed to be identical with the internal phenomenon itself. Someone can be in the same epistemic situation as he would be if there were heat, even in the absence of heat, simply by feeling the sensation of heat; and even in the presence of heat, he can have the same evidence as he would have in the absence of heat simply by lacking the sensation *S*. No such possibility exists in the case of pain and other mental phenomena. To be in the same epistemic situation that would obtain if one had a pain *is* to have a pain; to be in the same epistemic situation that would obtain in the absence of pain *is* not to have a pain. The apparent contingency of the connection between the mental state and the corresponding brain state thus cannot be explained by some sort of qualitative analogue as in the case of heat (151-152).

Pain is here just an example of a mental state. Kripke holds that the pain-brain state identity could be saved if when we try to imagine pain without the relevant brain state, we are merely imagining the *sensation* of pain existing without the corresponding brain state—and *not the pain* existing without that state. Other known identities will, after all, appear contingent if we conflate the sensation with the sensed phenomenon in our imagination. When we attempt, for instance, to imagine heat without molecular motion, we may find that we are actually just imagining the

[53] I take a quality space theory of pain to be the best rendering of the idea that each token pain is best characterized as a sensation *despite* there being no distinct pain quale. This view is taken up briefly at the end of this section and in more detail in section 3 with perceptual views.

sensation of heat without molecular motion—leaving the heat-molecular motion identity in tact. Kripke claims that no such salvation is possible for the pain-brain state identity, however, since pain just *is* a sensation.

Kripke later makes even more explicit his position that pain is essentially and exhaustively a sensation; i.e., pain just is its "immediate phenomenological quality," or what I have been calling its qualitative content. Consider:

> Pain… is not picked out by one of its accidental properties; rather it is picked out by *the property of being pain itself, by its immediate phenomenological quality.* Thus pain, unlike heat, is not only rigidly designated by 'pain' but the reference of the designator is determined by an essential property of the referent. Thus it is not possible to say that although pain is necessarily identical with a certain physical state, a certain phenomenon can be picked out in the same way we pick out pain without being correlated with that physical state. If any phenomenon is picked out in exactly the same way that we pick out pain, then that phenomenon *is* pain (152-153, emphasis added).

I take this passage to state that pain is essentially and exhaustively an immediate phenomenological quality.[54] Kripke's goal here is to use this claim to attack mind-brain identities generally. What exactly, however, is his argument for the claimed nature of pain? It is difficult to construct one. He seems to take it as intuitively obvious throughout the lectures, at one point rhetorically asking: "Can any case of essence be more obvious…."[55] Accordingly, we are regularly invited to consult our intuitions, and that would seem to be the end of the matter. Evaluating the claim that pain is a sensation, however, requires getting clear on what the intuitively-obvious qualitative content of pain is supposed to be.

---

[54] I am thankful to Jesse Prinz and David Pereplyotchik for pressing me on this interpretive point. It is admittedly textually unclear whether Kripke takes pain's qualitative content to be *exhaustive* of it in addition to being essential. If successful, the arguments here will show that there is no distinctive qualitative content either essential to or exhaustive of pain—whether Kripke is committed to both or only the weaker claim.

[55] Kripke, *Naming and Necessity*, 146.

Pains, *prima facie*, vary widely in their qualitative content; consider the qualitative difference (the difference in what it is like or how it feels) between stepping on a nail and a stomachache. The McGill Pain Questionnaire is, at present, the most widely used and respected tool for evaluating the contents of patients' token pains for diagnostic purposes. Ten major groups of common descriptors of the sensory aspect of pain are identified:[56]

> Group 1: Flickering, Quivering, Pulsing, Throbbing, Beating, Pounding.
> Group 2: Jumping, Flashing, Shooting.
> Group 3: Pricking, Boring, Drilling, Stabbing, Lancinating.
> Group 4: Sharp, Cutting, Lacerating.
> Group 5: Pinching, Pressing, Gnawing, Cramping, Crushing.
> Group 6: Tugging, Pulling, Wrenching.
> Group 7: Hot, Burning, Scalding, Searing.
> Group 8: Tingling, Itchy, Smarting, Stinging.
> Group 9: Dull, Sore, Hurting, Aching, Heavy.
> Group 10: Tender, Taut, Rasping, Splitting.

Which quality or combinations of qualities from the above list is essential to pain? It seems clear that none of these qualities can be the sort of thing that Kripke has in mind; whichever quality we choose, some pains occur with a different quality. On the idea that there is some such particular, distinct quale for pain, Gustafson (2005) bluntly says:

> There is no general way pains all feel. If this is so, we should wonder what defenders of it have in mind in the view that 'qualia' define what kind of state or condition a conscious experience is, essentially, and that, there must be a single 'pain quale' in virtue of which a state or condition is a pain, necessarily. What is the one feature of a feeling-sort that all and only pains have? Is the idea of such a thing credible? I think not. (219)

The obvious variation in qualitative contents suggests interpreting Kripke as claiming not that the *type* pain always has a distinctive qualitative content, but that every *token* pain has a

[56] Note that the questionnaire also provides descriptors of the affective and evaluative aspects of pain, and three miscellaneous groups of descriptor words that are not happily construed as sensory, affective or evaluative. It is worth nothing that this widely used questionnaire takes multiple aspects of pain to be components. I focus in this section on the variety of descriptors characterizing the sensory aspect alone, revisit the affective descriptors in section 4, and consider the questionnaire in its entirety in Chapter 4.

qualitative content essential to, but only to, that very token. This, however, does not work either. It is easy to imagine any particular pain being less burning, or stinging. So too, we at least *seem* to commonly experience the qualitative change of numerically identical pains. It is common to talk about our pains this way; e.g., *the* headache *was* throbbing and *is now* a dull ache, *the* pain in my stomach *was* sore and heavy, but *is now* jumping and lacerating. There is no independent reason for rejecting this way of talking. Most importantly, its diagnostic utility supports acceptance. If this useful way of talking is maintained, then the particular qualities identified by the McGill Pain Questionnaire cannot be the sort of qualitative content essential to *token* pains any more than the *type*. The distinctive qualitative content, understood as some combination of the above descriptors, is a contingent feature even of token pains.

Here it may be properly objected that none of the mentioned qualitative contents are the sort of distinctive, qualitative content Kripke had in mind. On the most natural reading, the intended qualitative content is something like the (much more general) *painfulness* of pain. By 'painful' I will mean the nasty, "this sucks" quality that is characteristic of many pain experiences. The *hurt*. The *unpleasantness*. On this reading, Kripke's claim that pains have an essential, distinctive, phenomenological quality is understood as the claim that pain just is a *painful* sensation; i.e., a mental state with a *painful* quale. The sharp pain of stepping on a nail and the dull pain of a stomachache, though varying in some aspects of their qualitative content, may both be *painful*. It is the painfulness that makes both experiences pain. On this reading, the Kripkean claim is challenged if we can imagine any token pain, without painfulness (i.e., hurt or unpleasantness).[57]

[57] Clark (2005) argues directly that painfulness could not be a quale. His argument, however, crucially relies on the plausible, but merely stipulated, claim that painfulness is a relation to desire or volition (185). His argument is explicitly aimed at Kripke and views like his (189). He

Fortunately, to evaluate this claim we need not wade through modality-deep metaphysical waters or exercise imaginative abilities. The refutation of the claim that pains are essentially painful lies in the *actual*. The intuition that pain just *is* a painful sensation may be deeply embedded, however, so it will be beneficial to consider a number of actual counterexamples and then to explain *away* the intuition itself.

One class of pains lacking painfulness is the chronic pains experienced by lobotomized patients.[58] As Nikola Grahek (2007) points out:

> The postoperative condition of patients who have undergone prefrontal lobotomies to treat chronic pain is often regarded as the paradigm case of an extreme reactive dissociation syndrome—or at least of striking indifference to pain. After the operation, these patients typically report that their pain is still present, but that it does not bother them anymore, and/or that they do not mind or care about it anymore (31).

The first-person reports in numerous case studies of such patients support the interpretation that these patients experience *at least some* of their pains without experiencing a painful sensation. Consider the following case study presented in Brand and Yancey's (1997) classic:

> By all measures, the lobotomy was a great success. The woman emerged from surgery completely free of the suffering that had shadowed her for a decade. … More than a year later I visited this couple in Bombay. The husband spoke enthusiastically about the lobotomy, and the woman herself seemed calm and content. When I inquired about the pain, she said, "Oh yes, it is still there. I just do not worry about it anymore." She smiled sweetly and chuckled to herself. "In fact, it is still agonizing. But I do not mind."
>
> At the time it startled me to hear words about agony coming from a person with such a placid demeanor: no grimace, no groan, only a gentle smile. As I read

---

says: "… unless a sensory episode stands in the appropriate relations to one's motivational states, it is not aversive, and hence not painful. Hence painfulness is not a quale" (187). I do not follow Clark in giving this argument here, since I find it question-begging in the current context to take painfulness to be a (functional) relation to one's motivational states—if painfulness is the quale that Kripke has in mind, this will not simply be accepted.

[58] In lobotomy, connections to prefrontal cortex are severed by incisions in the frontal lobe.

> about other lobotomies, however, I found she was displaying a very typical attitude. …
>
> Lobotomized patients rarely ask for medication. As a German neurosurgeon who had performed many frontal lobotomies once told me, "The procedure takes all the suffering out of pain," (210-211).

In *Psychosurgery in the Treatment of Mental Disorders and Intractable Pain*, Freeman and Watts (1950) present a number of case studies describing the experiences of lobotomized patients. Consider this choice excerpt:

> Case 301… Pain in back since girlhood, worse since birth of first child 28 years ago, and unbearable for last three years. … Prefrontal lobotomy of the standard type… She says the pain is still there, but it does not bother her (358).

It is important to note that only "at least some" pains are experienced by lobotomized patients as non-painful; preexisting chronic pains are often no longer painful, but fresh, acute pains often still are. So too, the painfulness of the preexisting pains may sometimes simply vary in degree, though not in kind, from the original. As Melzack and Wall (2008) describe lobotomized patients:

> Typically these (frontal lobotomy) patients report after the operation that they still have pain but it does not bother them. When they are questioned more closely, they frequently say that they still have the 'little' pain, but the 'big' pain, the suffering, the anguish are gone. … [T]hese patients may complain vociferously about pinprick and mild burn. … The aversive quality of the pain and the drive to seek pain relief both appear to be diminished (137).

If we take patients' reports and descriptions by doctors and pain researchers of the lobotomized at face value, we can conclude that there are at least some cases of pains that are not painful.

Though I think interpreting lobotomized patients' chronic pains as non-painful is the most natural, it is admittedly not the only interpretive option. Two others should be considered. First, one may say instead that the pains of lobotomized patients *are* painful, but that the patients are simply untroubled by the painfulness. Brand and Yancey's description of the woman who

still reports her pain as "agonizing" may suggest this interpretation. This alternative, however, does not do justice to the seemingly very different postsurgical *qualitative* content. As Melzack and Wall point out, as a result of a lobotomy the aversive quality of the chronic pains are removed—that very quality we are currently interpreting Kripke to be claiming was essential. Second, one may say, as suggested by Melzack and Wall, that the painfulness of the lobotomized chronic pains are still always present, but are either diminished or not consciously experienced. Note that this interpretation, however, even if granted, still suggests that the painfulness of pain is an independently varying *component* of pain. Even on this interpretation, the painful qualitative content of pains is not *exhaustive* of pain; pains are not *simply* a painful sensation. Moreover, if this component independently varies, we should at least be able to *imagine* it being completely abolished—so it is not essential by Kripke's criterion. Most important, there is no independent reason to maintain that the painfulness of pain just *must* be non-consciously present. For those unsatisfied that lobotomized patients present examples of pains that are not painful, however, we turn to pain asymbolics.

The history of the term 'pain asymbolia' leaves plenty of room for disagreement about its extension and how its possessors are best described. There are a number of pain dissociation syndromes and the terms used to refer to them have not always been employed systematically.[59]

As defined in the classic paper, "Asymbolia for Pain," by Rubins and Friedman (1948):

> Asymbolia for pain denotes the inability to recognize the unpleasant or disagreeable component of a painful or threatening stimulus, with the result that little or no defense reaction is produced, although the noxious stimulus itself is perceived (554).

[59] See Nagasako, et al (2003) for a review of the history of this and related terms.

Note that the unpleasant or disagreeable component (i.e., painfulness) of pain is here described as a *component* of pain that asymbolics are unable to recognize. The definition stays neutral concerning whether the painfulness of pain is not present, or present but not perceived.

The stated criteria for the diagnosis of pain asymbolia employed by Berthier, et al. (1988) in their classic study of 6 patients is similarly opaque. In addition to requiring that the subject be of sound mind, to be diagnosed pain asymbolic it was required only that subjects "… had no deficit in pain perception; and had absent or inadequate motor or emotional responses to painful stimuli applied anywhere on the bodily surface."[60] This criterion stays neutral concerning whether the painfulness of pain is entirely absent, or present but not effective for issuing in behavior—in addition, of course, to the loose disjunctive requirement that *either* motor *or* emotional responses be absent.

Both classic definitions of pain asymbolia, then, leave open the possibility that the painfulness is still present, but merely not conscious, not perceived, or not causally effective. Note, however, that even these interpretations allow us to conclude that painless pains are imaginable. Not *aware* of the painfulness (even if it is there), pain asymbolics report pains that are not painful. Such patients not only imagine pains that are not painful, they report that they are *experiencing* them. The claim that painfulness is essential to pain is thereby undermined if we rely on imagination and possibility as Kripke does, but there is a more direct objection. However loosely we *define* the condition of pain asymbolia, there is at least a subset of individuals diagnosed with it that are accurately described as experiencing pains that are not

[60] Berthier, "Asymbolia for pain: A sensory-limbic disconnection syndrome,"42.

painful. I take this to be now well established within the pain research community and consider just a few classic cases here.[61]

Despite their rather loose and (for our purposes) equivocally stated criteria, Berthier, et al. (1988) determined that their patients had pain experiences that were not experienced as at all unpleasant:

> Although all 6 patients could adequately recognize painful stimuli and distinguish sharp from dull, all of them showed a lack of response to painful stimuli applied over the entire body. Neither superficial nor deep pain stimulation elicited a motor withdrawal, grimacing, or an appropriate emotional response. … Inappropriate emotional responses were common: 4 patients smiled or laughed during the pain testing procedure. … All patients appeared quite unaware of their abnormal reactions… None of them became angry or anxious during the pain testing procedure; *in fact, while all could recognize pain, none of them reported any unpleasant feeling.* Patients showed normal autonomic reactions (tachycardia, hypertension, sweating, mydriasis) during the painful stimulation, but failed to react with a flinch, blink, or adequate emotional responses… (43, emphasis mine).

These patients at least seem to be experiencing pains that are not painful, and their doctors describe them that way.

Consider the colorful case study of a teenage boy presented by Frances and Gale (1975):

> During pubescence he participated in sports and motorcycling with fearless abandon and an exhibitionistic flair. He began to demonstrate his 'freakish' but impressive insensitivity to pain by applying lighted cigarettes to various extremities. This often became a contest with other boys for both money and pride, with the patient the undefeated champion. He had no apprehension over the disfigurement and only later came to regret it because of a growing awareness that the scars made him less attractive to others, particularly girls. …In spite of his clear cognitive understanding of his vulnerability to injury and the experience of numerous broken limbs, the patient participated without restraint in contact sports and frequent fist fights. …The patient describes his injuries with an indifference which would seem to belie the fact that they had happened to him. Paradoxically, he would be more upset by damage to his clothing than by bodily injury (110).

[61] For more case studies see, for instance: Rubins and Friedman 1948; Berthier et al. 1988, 1990; Dong et al. 1994; Weinstein et al. 1995.

Unfortunately, as another testimony to the tortured use of pain terminology, these authors conflate what pain researchers now distinguish: pain insensitivity and pain indifference. Pain insensitivity is now agreed to be a condition whereby individuals do not experience *any* sensations *as* pain and *cannot* reliably discriminate among various noxious stimulations. Pain indifference, contrastingly, is a condition whereby individuals *can* reliably discriminate among noxious stimuli and experience some sensations *as* pain, but are indifferent to those experiences. The case in question is of an individual who reported pain, knew that he was being injured, but was not bothered by his pain. He would now be properly diagnosed as having pain indifference, not pain insensitivity; he feels pain, but the pain he feels is not experienced as painful.

Confronted with these cases of pain asymbolia, retaining the claim that the painfulness of pain is essential to it requires embracing one of two options: First, that what these individuals are experiencing is simply not *pain*, or, alternatively, that they are wrong about whether the pain they are experiencing is painful. Both options are desperate.

The lack of withdrawal response or aversive reaction one would predict of a *painful* experience tells against declaring that the sensations are painful and the subjects are wrong about it. Indeed, independent of the essentiality intuition in question, it is hard to see *any* reason, first or third-personal, for construing these subjects as undergoing painful sensations—including non-consciously.

Denying that what the subjects are experiencing is pain *at all* is much more tempting. Indeed, some even in the medical profession succumb to this temptation. Consider this heartrending case described Osuntokun, et al (1968):

> … a 5-year-old boy … rarely cried, but shed tears normally when he did and appeared incapable of being easily hurt, his parents describing him as ‘a very strong boy’. He was very prone to ulcers on his feet. … There was no history of excessive or diminished sweating [a typical symptom of pain insensitivity as

> against pain indifference] … Neurological examination showed normal sensorium. … There was no detectable abnormality in the motor system, the muscles, the gait, and coordination being normal. In the sensory system light touch was normal. … His reaction to noxious stimuli was, however, odd. He knew when he was being pricked by a pin, for he could tell the difference by gesture between the sharp and blunt ends. He showed no withdrawal response, evasive or preventive motor reactions, and no facial expression of pain to pin-prick venipuncture, intense pressure on the Achilles tendon, or to his skin being transfixed with a needle—which he seemed to enjoy very much. He invited the examiner to repeat the tests while he remained all smiles. … *The patient was in hospital for eight weeks, during which news got around of the 'wonder boy' who could not feel pain*. Subsequently … some laboratory technicians surreptitiously paid him several visits, transfixing him with pins and needles and watching with awe how much he enjoyed it. He appeared, however, to discover that he was being regarded as a 'freak' and developed the greatest dislike of pins and needles. He would fly into an uncontrollable rage and temper tantrums at the slightest exhibition of a pin or needle, but would allow himself with much pleasure to be subjected to other forms of noxious stimuli like pressure on the tendon Achilles and hot water at 75° C (291-292, emphasis added).

The ability of this child to discriminate among noxious stimulation tells against a diagnosis of pain insensitivity, instead supporting one of congenital pain indifference—as the doctors in this case indeed diagnosed him. Notice, however, that rumors of the subject spread that he could not feel pain *at all*.

Two major theoretical considerations point against interpreting pain asymbolics, like this child, as not experiencing pain.[62] First, the individuals report their sensations as pains, and unless we have good reasons to think otherwise, first-person reports should be accepted as veridical. Indeed, as we will see in the next chapter, the International Association for the Study of Pain goes too far in making first-person reports the *only* appropriate criterion for whether an individual is in pain. One might think that the first-person report actually provides *support* for the sensation account of pain; surely it is in virtue of their qualitative content that these patients

[62]A further consideration not discussed here is whether there is activation of the same neural mechanisms in these patients as those correlated with the sensory component of paradigmatic pain experiences in individuals who are not asymbolic.

identify their experiences as pain. Note, however, that if their identification *is* in virtue of some qualitative content of the experience, *they* do not think the relevant qualitative content is *painfulness*, since *they* do not think that their pains are *painful*. Second, and more important, are individuals' abilities to discriminate the intensity and other qualities of their pain sensations. If we are inclined to doubt the veracity of their reports, their discriminatory abilities speak in favor of accepting them. If individuals report their experiences as pains, and can reliably discriminate among the qualities and intensities of their pains like non-asymbolics, insisting that what they are experiencing is not *really* pain would appear to be arbitrary linguistic legislation.

There seems to be no reason, independent of the essentiality claim in question, that can be offered for denying that these cases are cases of pain. We ought, instead, to conclude with Grahek that, "Although … pain without any painfulness may seem inconceivable or incomprehensible, we must concede that it is possible, given that abundant clinical evidence attests to its existence."[63] Pain asymbolics constitute a refutation the claim that the painfulness is essential to pain; not only can we *imagine* cases of pain without painfulness, pain asymbolics constitute *actual* cases.

I can think of only one more construal of the immediate phenomenological quality held to be essential to pain. Perhaps the distinctive pain quale is actually *any* quale that is a member of a circumscribed set. For reasons to be developed in the next section, call this set a quality space. This view allows token pains to vary in qualitative content, while nonetheless characterizing the type pain by qualitative content. No one point within the quality space is taken to be essential to a token mental state being a pain. Rather, what is essential is that the mental state has a qualitative content characterized as lying somewhere within the pain quality space.

---

[63] Grahek, *Feeling Pain and Being in Pain*, 50.

This construal of Kripke's essentiality claim is, however, an unhappy one. If we construe him as claiming that pains essentially have a qualitative content lying somewhere in a pain quality space we are simply no longer committed to the idea that pain *is* a particular quale. Pains, on this view, would *vary* in qualitative content, even as visual sensations vary in qualitative content. Accordingly, I interpret this view to be that pain is a distinctive modality—not just a particular quale.[64] I turn to the refutation of perceptual views of pain, including the view that pain is a distinct modality characterized by its proprietary quality space, in the next section.

Before turning, however, the strong *intuition* that pains are essentially painful deserves further comment. Despite the above arguments, this intuition may be hard to shake. For those of us who have never (at least not consciously) experienced pains that are not painful, the very idea of them may remain dubitable. I think the strength of this intuition can be explained, and so hopefully explained *away*, by acknowledging that what is most *important* to us about our pains is their painfulness. Pains that are not painful are not bothersome. Painful pains, on the other hand, can be so bothersome that they obliterate our ability to concentrate on anything else.

It seems to me that those, like Kripke, who hold that pain is essentially a quale—in particular, that it is essentially *painful*—have made an understandable mistake. They have confused what is most *important* to us about pain, at least while it is occurring, with what is *essential* to pain. This mistake is one that Kripke himself (1980) warns against:

> Theorists have often said that we identify objects across possible worlds as objects resembling the given one in the most important respects. On the contrary, Nixon, had he decided to act otherwise, might have avoided politics like the plague, though privately harboring radical opinions (53).

[64] This is as it should be for quality-space theory, since according to these views, families of mental qualities (sensations) are characterized by the perceptual discriminations they enable within a modality.

It is mistaken to think that Nixon was essentially involved in politics. It is admitted, however, that it tempting to say that Nixon wouldn't have been Nixon if he had not engaged in politics. This temptation is explained, and hopefully explained *away,* by pointing out that it confuses essentiality with importance.

Again, and as a general matter:

> *Important* properties of an object need not be essential, unless 'importance' is used as a synonym for essence; and an object could have had properties very different from its most striking actual properties, or from the properties we use to identify it (77).

Applying these ideas to pain: The importance of the painfulness of pain may lead one to infer that painfulness is therefore essential to it. As Kripke himself argues, this inference is illegitimate. The temptation to think of painfulness as essential to pain, despite actual cases of pains that are not painful, is neatly explained by our confusing essentiality with importance. The intuition that painfulness is essential to pain may be very strong, but actual cases of pains that are not painful warrant our abandoning this intuition.

Pain is not simply a sensation. There is no distinctive qualitative content essential to all pains. Indeed, for any particular quale we might have in mind, not only can we imagine pains without it, there are *actual* cases of pains without it — this includes painfulness.[65] While pains often have qualitative content, characterizing pains solely in terms of qualitative content will rule out states that we have independent reasons to think are pain. There is, that is, also no qualitative content exhaustive of pains. I have focused on Kripke's arguments in this section, but the refutation of the unitary account of pain as a sensation generalizes. There is no qualitative content essential to or exhaustive of pain; pain is not simply a sensation.

---

[65] The essentiality of painfulness when it is characterized as an emotion instead of a quale is discussed in section 4.

3. Pain is Not Simply a Perception

In this section, I consider the view that pain is simply a perception. Two of the most influential proponents of this popular view are David Armstrong (1962) and Michael Tye (2000, 2006). Both views characterize pains as perceptions garnered through the modalities of touch or bodily perception. I refute each in turn before considering whether a perceptual view that characterizes pain as a distinct modality fares better. The best candidate for the thing perceived is tissue damage. The overarching problem with perceptual views, however, is that though pains often involve a perception of tissue damage, characterizing pain *solely* as such a perception illegitimately deems too many pains hallucinatory or illusory and characterizes too many states perceptions as pain that are not. While positing excessive misperceptions is not alone decisive, these "misperceptions" are better explained by characterizing pain as an experience with multiple components. Perception of tissue damage is neither necessary nor sufficient for pain. Considering pain to be a modality of its own with a distinctive quality space fares no better: there are no good candidates for the dimensions of a distinctive space, and pain lacks all of the paradigmatic characteristics of a perceptual modality. Pain is simply *not* simply a perception.

According to Armstrong, perception is a process whereby an organism acquires beliefs, and beliefs are characterized by the discriminatory behavior they allow. Armstrong distinguishes two kinds of perceptions: mediate and immediate. Immediate perceptions are perceptions of what Armstrong calls the *sensible qualities* of objects. Each sensory modality has a distinctive, proprietary class of sensible qualities immediately perceived through that modality. Mediate perceptions, by contrast, are perceptions of *objects* that are automatically suggested or conveyed by immediate perceptions. Consider Armstrong's example of perceiving a cat: I immediately perceive the object's visual sensible qualities—identified as color, light and shade, shape, size,

motion, and spatial relations. The *cat*, however, I perceive mediately in a "perfectly automatic" way through what is suggested or conveyed to me by my immediate perceptions.[66] Again, the core of the distinction is that mediate perceptions are perceptions of objects and immediate perceptions are perceptions of sensible qualities. Importantly, objects and sensible qualities are both taken to enjoy an objective existence, and are thereby appropriately contrasted with misperception and subject to accuracy conditions.

Central to his account of pain is Armstrong's further distinction between immediate perceptions and what we only *seem* or *appear* to immediately perceive. These appearances or seemings are called sense-impressions and characterized as abstractions from immediate perceptions. What we are taken to abstract away *from* is the belief that the sensible qualities are as we perceive them to be. I may, for example, have a green sense-impression upon perceiving something I antecedently believe to be blue. In such a case, I may persist in believing that some mediately perceived *object* — a book cover, say — is blue, while nonetheless undergoing a green sense-impression. This is intended to accord with our common ways of reporting perceptions and sense impressions: The book cover *looks* green, but I nonetheless *believe* that it is blue

On Armstrong's view, all bodily sensations, including pains, are best characterized not as full-fledged perceptions, but as bodily sense-impressions. According to Armstrong, bodily sense-impressions are abstractions from what is immediately perceived through tactition—whereby we perceive the relation between our bodies and external objects—or bodily perception—whereby we perceive our bodies themselves. Just as we use the familiar phenomenological sense of '*looks'* when abstracting from visual perceptions, we often use a phenomenological sense of the word

[66]On my reading, 'automatic' means only *seemingly* non-inferential.

'feels' to report or express the sense-impressions abstracted from bodily perception or tactition.[67] Using 'feels' this way: It may *feel* as if my hand is hot, though it is not and I do not believe it is.

Some bodily sensations *prima facie* resist an appearance/reality distinction, and Armstrong calls these intransitive. Among such sensations Armstrong includes tickles, aches, itches, erotic sensations, and, of course, pains. According to Armstrong, we *seem* unable to distinguish between having an itch and the sensation of an itch, and between pain and the sensation of a pain. As bodily sensations are themselves characterized as seemings, the lack of a distinction between the sensation and the thing sensed for intransitive sensations amounts to a lack of a distinction between sensing a thing and *seeming* to sense it. In such cases, it seems impossible to distinguish between the relevant immediate perceptions and the abstracted sense-impressions. For the transitive sensations, such a distinction is quite intelligible; we can intelligibly distinguish between sensing the heat of our hand and seeming to have a hot hand. For the intransitive sensations, however, there *seems* to be no distinction between seeming to have the sensation and having the sensation—no distinction between having a pain and seeming to have a pain.

It is important to be clear, however, that Armstrong's view is *not* that intransitive sensations are pure sensations; they have intentional content. Consistent with their status as mere sense-impressions, this content is an abstraction from any *belief* in the objective existence of the relevant sensible qualities. As such, the intentional content of pain states are characterized as *as*

[67] On Armstrong's view, tactition and bodily perception should not be understood as rigidly distinct because all tactition also involves bodily perception. The difference is that there is an organ of perception for tactition and none for bodily perception. For more on the relation between tactition and bodily perception, see Chapter 4, of *Bodily Sensations*. The problems I raise for Armstrong's perceptual view arise regardless of whether pain is characterized as an abstraction from tactition, bodily perception, or both.

if statements. To be in pain is to feel *as if* something is occurring, but it is not to *believe that* something is occurring.

The "as if" construal of the intentional content of sense-impressions allows Armstrong the appearance/reality distinction that he thinks *prima facie* impossible to apply to the intransitive sensations. To have a sense-impression is to feel *as if* something is occurring, but we may sensibly go on to ask if what it feels *as if* is occurring is *actually* occurring. Importantly, however, even if the answer is no, we are still having the sense-impressions: we are still having a pain even if what it feels *as if* occurring is not occurring. Indeed, even if we do not believe that what it feels as if is occurring is actually occurring. This is part of the intended utility of *abstracting* from the beliefs that are crucial to characterizing the perception.

In the case of pain, what does it feel *as if* is occurring? Armstrong's answer: Bodily disturbance at, or in, a location. Bodily sensations are held to be especially indistinct, such that we may be imprecise about the nature and location of the disturbance. To be a pain, however, it must feel as if there is *some* sort of disturbance occurring *somewhere* in the body.[68] Since pains are not full-fledged perceptions, we may not *believe* there is any disturbance at a bodily location, but to be pain it must *feel as if* there is.

So far, this account ignores painfulness. As discussed in the last section, painfulness is so important that it is tempting to think it is essential, or even exhaustive, of pain. Armstrong avoids these mistakes. Pains (bodily sense-impressions as of disturbances at locations) are held to characteristically *cause* a distinctive reactive attitude dubbed "the pain reaction." Painfulness, he maintains, is the relation of the sense-impression to that attitude. Our everyday *concept* of pain

[68] Armstrong distinguishes between bodily feeling and bodily sensation on the basis of whether the sense-impression has an intentional location. Pain is a bodily sensation and not a bodily feeling since it is necessarily has an intentional location. See Chapter 8 of *Bodily Sensations*.

includes both the bodily sense-impression and the characteristic reaction to that impression, but Armstrong maintains that the causal relationship between the two is contingent. Quoting at some length:

> Having a certain sense-impression does not entail taking up any particular attitude to that impression. It is a contingent fact that we like or dislike certain bodily sense-impressions. But it is now being suggested that the concept of physical pain is a portmanteau-concept: that it involves *both* the having of a certain sort of bodily sense-impression, *and* the taking up of a certain attitude to the impression. The *painfulness* of the sense-impression is therefore a *relational* property of the impression, in just the same way that the painfulness of a scene is a relational property of that scene. We call physical pain 'pain' simply because the having of these impressions of bodily disturbance is normally the most painful thing in human life (107, *Bodily Sensations*).

Painfulness, on this view, is a relation between the sense-impression that constitutes pain, and the pain-reaction that the sense-impression characteristically causes. The pain-reaction is taken by Armstrong to include an immediate (seemingly unmediated) and interested (the reaction is to *my*, and only *my*, sensation) dislike of the feeling, and a concern for the disturbed location. It *sounds* unnaturally redundant to speak of painful pains because pains normally cause the pain-reaction, but on Armstrong's view a pain that is not painful is neither logically contradictory nor nomologically impossible.

The contingency of the pain-reaction allows Armstrong to explain the cases that proved intractable for the central interpretation of the Kripkean view. Pains that are not painful occur when the pain-reaction is, for any reason, not caused by the sense-impression. Armstrong offers the asymbolias and lobotomies discussed in the previous section, along with mild pains, as examples of pains that do not cause the normal pain-reaction. In these cases, "… a linguistic decision has to be taken as to whether they can be said to be feeling 'pain' or not."[69] We may, that is, not wish to *call* such cases pain. Note, however, that the linguistic decision here will not

[69]Armstrong, Bodily Sensations, 108.

alter which sense-impression one is having, however, and Armstrong's view is that pain is an intransitive bodily sense-impression.

Armstrong's view about pain is, then, in sum that:

> To have a pain in a certain place … is to feel a disturbance of our normal bodily state at that place; together with an immediate and interested dislike of that feeling; and a concern for the place where the disturbance feels to be. … Physical pain becomes a certain sort of highly indeterminate bodily sense-impression, an impression that happens to arouse certain (quite determinate) attitudes in the person who has it (105, *Bodily Sensations*).

Despite Armstrong's beautifully subtle account, his view ultimately fails for the same reason as any other perceptual account. Perceptions the account deems *mis*perceptions and deems pains that are *not* pain, are both so prevalent that it is not credible to think of them simply as abnormal workings of the relevant modalities. The profound lack of correlation between located bodily disturbance and pain is *so* profound that while perception should be embraced as a typical *component* of pain, it is implausibly that it is the whole story.

Recall that on Armstrong's view pains are not *full-blown* perceptions—that is, they do not commit the bearer to any *beliefs* about whether there is a located bodily disturbance. In standard parlance: The locations are merely *intentional* locations and the bodily disturbance is a merely *intentional* bodily disturbance. The "mere" intentionality is captured by the "as if" construal of pains' contents. This suggests that the correlation—or lack thereof I argue for below—between bodily disturbance and pain is irrelevant.

Against this, however, recall also that Armstrong's sense-impressions still admit of an appearance/reality distinction such that it is sensible to ask whether they are representing the world as it actually is. This makes sense since, on Armstrong's view, *even if pains are merely sense-impressions, there must be some immediate perceptions from which we abstract to characterize them.* Pains are held to be abstractions (we drop the belief) from the immediate

perceptions of tactition and bodily perception. As these perceptions are to be contrasted with misperceptions, the correlation between bodily-disturbance-at-a-location and pain is actually quite germane. Excessive mismatch between bodily-disturbances-at-a-location and pain would tell against characterizing pains as (abstractions from) perceptions that there is such a disturbed location.

Armstrong himself sees the importance for his account that pains are at least usually accurate in their representation of bodily disturbances at bodily locations. We must usually be right about both *whether* and *where* a bodily disturbance is happening. By Armstrong's lights, if we are incorrect about *where* the bodily disturbance is we are undergoing an illusion, and if we are wrong about *whether* there is bodily disturbance at all, then we are undergoing a hallucination.[70] Hallucinations and illusions, however, should be exceptions to the otherwise accurate representations of the relevant modality, and Armstrong recognizes this without finding it a problem. He thinks that pains usually *do* accurately represent bodily disturbances at bodily locations. When the sense-impressions *do not* accurately represent, they will still be pains—that is the utility of characterizing them as sense-impressions instead of full-blown perceptions. Pains that inaccurately represent, however, will be abstractions from *misperceptions*—that is why they will also be illusions or hallucinations. On Armstrong's view then, a hallucinatory or illusory pain is a real pain, but it is also a real hallucination or illusion.

There are a great many cases that Armstrong's, or most *any* perceptual view, deem hallucinations, illusions, or misperceptions. The fact that there are *so* many tells strongly against characterizing pains as (abstractions from) perceptions of disturbed bodily locations.

---

[70] Armstrong, *Bodily Sensations*, (55-57, 108-109).

A brief terminological detour is necessary before turning to a brief catalogue of cases. Notice that feelings as of bodily disturbances at locations will not just yield sense-impressions that are *pain*. The intentional content of itches, erotic sensations, tickles, and skin irritation could all be equally well characterized as sense-impressions as of bodily disturbances at bodily locations—though each have different, contingent, characteristic reactive attitudes. This may itself be thought a serious objection to Armstrong's view. Abstracting from the immediate perception of being tickled, for instance, will yield a content as if of a bodily disturbance in my foot, but it will *not* (usually) yield pain.[71] To avoid this problem, perceptualists traditionally and most commonly construe the contents as being (as) of *tissue damage*. This does indeed seem the best candidate for the type of bodily disturbance represented in pain. The most charitable thing to do seems to be to focus on dissociations between tissue damage and pain, since were Armstrong's general construal used, the dissociations would be yet more plentiful.

Let us now consider some of the many types of pain, a number of them quite common, that Armstrong's account classifies as hallucinations or illusions.

I. Headaches: Headaches are extremely common, but involve no known tissue damage—certainly none to tissues in the head where the pain feels to be.

II. Lower-back pain: Another common condition, lower-back pain usually involves no damage to the lower back. Loeser (1994) reports that there is no known bodily damage in 70% of these cases.

[71]The worry is not only that many other bodily sensations could be characterized with the same intentional content, but that ruling these out seems to require including qualitative or affective content distinctive of pain as components. Armstrong discusses this sort of worry in *Bodily Sensations* (113-119).

III. Pain after healing: People undergoing this robust phenomenon have healed injuries followed by increasing pain. Discussing these sorts of cases as a "syndrome," Livingston (1943) says:

> The onset of symptoms may follow the most commonplace of injuries. A bruise, a superficial cut, the prick of a thorn or a broken chicken-bone, a sprain or even a postoperative scar may act as the causative lesion. The event which precipitates the syndrome may appear both to the patient and the physician as of minor consequence, and both have every reason to anticipate the same prompt recovery that follows similar injuries. This anticipation is not realized and the symptoms tend to become progressively worse (110).

IV. Referred Pain: There are *many* kinds of referred pain, and most are still poorly understood. The most well known is the referred pain of myocardial infarction; during a heart attack, pain is referred to the left hand or arm. A tightness, but no pain, is felt at the place of actual damage in the heart.

V. Phantom Limb Pain: Those unfortunate enough to lose a limb are likely to experience pain as if in the missing limb despite its absence. Though estimates vary, approximately 80% of amputees experience phantom limb pain.

VI. Thermal Grill Illusion: A grill is comprised of alternating warm and cool bars. When both kinds of bars are touched simultaneously, the experience is as of burning pain in the hand.

VII. Chronic Pain: Chronic pain is alarmingly prevalent and proliferating. It is most broadly characterized as pains either persisting beyond, or originating from something other than, known tissue damage. Chronic pain conditions include, but are by no means limited to, fibromyalgia, irritable bowel syndrome, sympathetically maintained pain, chronic fatigue syndrome, rheumatoid diseases,

and chronic complex regional pain syndrome. In a poignant description of the problem, Hardcastle reports:

> Chronic pain is an enormous and costly industry. Somewhere between 11 and 34 percent of all adults in the Western world suffer from some form of debilitating and ongoing pain at any given time. Persistent pain costs the United States somewhere between $40 and $100 billion a year in medical services, loss of productivity, and compensation payments. It is the primary reason for missing work and the primary cause of disability. It is the second most frequent illness—the common cold is the first—and affects about four-fifths of all people. It is the most prevalent symptom in patients seeking medical assistance and the primary motivation in seeking a physician in 80 percent of all doctor visits (9-10, *The Myth of Pain*).

These are all cases of pain without tissue damage either at all or at the intentional location. On Armstrong's account, they are all, therefore, cases of hallucination or illusion. This may seem acceptable for some cases—like phantom limb pain or the thermal grill illusion—but it is less acceptable for others—like headaches or chronic pain. Are headaches really best characterized as hallucinations?

In addition to the many pains without the perception of located tissue damage, there are also many perceptions of tissue damage with no pain. Armstrong himself considers two kinds of such cases: unconscious pains and unattended pains. His accounting of these dissociative cases faces its own problems of adequacy, which I put aside.[72] More problematic for Armstrong, however, are attended and conscious (or, at least, reported) perceptions of bodily damage, abstracting (the belief component) from which does not yield pain. Armstrong appears not to discuss any such cases, but consider the following:

I. Episodic Analgesia: Quite common, episodic analgesia is an isolated or transient experience of tissue damage unaccompanied by pain. Injuries incurred in battle or

[72] Armstrong considers these cases to be *unfulfilled possibilities* of sense-impressions (48-51). That we can account for an unattended or unconscious pain solely as an unfulfilled possibility strikes me as highly dubious, but I will not argue for that claim here.

sports are well known, but there are *many* types. Melzack et al (1982), for instance, conducted a study that concluded that thirty-seven percent of arrivals to the ER at the considered hospital reported that they did not feel any pain until minutes, or even hours, after the injury. Note that these cases differ diagnostically from shock—which may itself be considered a dissociative phenomenon.

II. Congenital Insensitivity to Pain: Those congenitally insensitive to pain do not report, and in general are considered unable to feel, any pain. Individuals suffering this condition almost always die very young, but a particularly famous case is Miss C who, remarkably, lived to be twenty-nine. Her meticulous, accurate observations of her own body seem to have been crucial contributors to her extended survival. She was quite able to perceive bodily disturbances and damage, but experienced no pain.

III. Indian hook-swinging ceremonies: In some Indian villages, a man is chosen to represent a god and bless both people and crops while swinging over them from above. To accomplish this, two large, steel hooks are attached by rope to a pulley-contraction and pushed through the skin and muscles of his back. The men undergoing the ritual neither seem to experience nor report any pain, but instead are happy and elated.

IV. Trepanation: Trepanation is a primarily East African practice to relieve chronic pain. The skull is exposed and scraped and no pain-medications of any kind are provided. The patients do not report any pain, sit calmly, and hold containers to catch the falling blood.[73]

[73]An American documentary, "A Hole in the Head," depicts the practice in vivid detail.

V. Couvadé: Couvadé a childbirth custom reported as occurring in numerous tribal cultures and more recently in New Guinea and Brazil. Delivering mothers undergo labor with only a few hours of minimal discomfort during labor before returning to work in the field. Expectant fathers, on the other hand, spend days on bed-rest in (at least reported and seeming) agony.

VI. Self-Injurious Behavior: Self-injurious behavior is most broadly characterized as any intentional, self-inflicted harm not intended to be fatal. "Cutting" is perhaps the most well known in American culture. Individuals engaging in SIB are well aware of the damage they are causing to their bodies, but often do not report or appear to feel any pain. On the contrary, SIB is often considered by clinicians, and is almost always reported to be, a *coping* mechanism that soothes or comforts the individual engaging in the practice.

These are all cases where bodily disturbances at locations are perceived (conscious, attended perceptions), and involve no pain. On Armstrong's view, these perceptions should yield pain if we abstract from the belief that things are to these individuals as they seem to be. In these cases, however, it does not *seem* to the individuals as if they are in pain any more than they *believe* that they are in pain—though it *does* seem to them (and they even believe) that there is tissue damage happening at a bodily location. Are all of these people really in pain though they do not report or act like they are? As before, this is more plausible for some cases than others. Should we really say that the boxer, for instance, is in pain and just does not know it? On the other hand, should we say that he does not perceive that he is undergoing tissue damage?

The moral of these stories is that the link between tissue damage and pain is tenuous. As we will see in the following two chapters, the poor correlation between pain and tissue damage is now expected and well-accepted by pain scientists.

I think it not implausible that if Armstrong had known the vast number of pains his account would classify as hallucinations or illusions and the many non-pains that would be classified as pains, he would have changed his view. Near the end of *Bodily Sensations*, he writes:

> Finally, as we have mentioned, illusions of bodily sense are, as a matter of fact, very infrequent. A visual illusion is a common thing, but the case of a 'phantom limb' excites wonder. This means that, in the vast majority of cases, the place where the bodily occurrence feels to be, and the place where it is, are identical. Given the place of the pain, we are given the place of the trouble. (It is a recent medical *discovery* that many internal pains are not at the site of the trouble. Such discoveries have not yet affected our ordinary concept of pain, although they might do so in the end.) (125).

The "end" that Armstrong speaks of in his parenthetical remark is here. I do not think that Armstrong's view ever adequately captured the ordinary concept of pain; consider headache and chronic pain on the one hand, and SIB on the other. It did and does, however, capture many *philosophers'* conception of pain. It is time to revise that concept. The mismatch between located tissue damage and pain is now known to be too prevalent for the perceptual view to remain credible.

Nevertheless, Michael Tye (2000), defends a contemporary perceptual view of pain as representing tissue damage and its qualities. Tye is a representationalist, holding that all content, both intentional and qualitative, is reducible without remainder to representational content. Tye's view differs from Armstrong's in a number of respects, and I will mention three.[74] The first is that Tye takes the affective or emotional aspect of pain to be represented in pain, whereas

[74] For a richer comparison of Tye and Armstrong, see Aydede (2009).

Armstrong takes it to be a typical *reaction* to pain. Second, while Armstrong takes pains to be sense-impressions, Tye takes them to be full-blown perceptions. These two differences both leave Tye even more vulnerable to the objection that tissue damage and pain are poorly correlated. The third way Tye's view differs crucially from Armstrong's, however, may ameliorate some of the problems raised by the other two. On Tye's view, not *all* perceptions (representations) of tissue damage are pains. Rather, the representation of tissue damage is only pain when it is represented *in a "pain experience."*

On Tye's view, the word 'pain' has two senses that are liable to cause confusion when conflated: the pain experience and tissue damage as represented in a pain experience. He holds that distinguishing these senses allows us to understand what we mean when we say both that the pain is *in* our leg ('pain' in the sense of tissue damage) and that *we* are in pain ('pain' in the sense of a pain experience). Note that this first sense can apply only if the second does; i.e., I cannot have a pain in my leg without undergoing a pain experience.

Three hints about what is required for an experience to be a *pain* experience crop up throughout Tye's discussion: tissue-damage in a location, badness, and a distinctive, "special something." This first possibility is a necessary but not sufficient condition. Without *some* requirements beyond the representation of tissue damage for when an experience is a *pain* experience, the view is vacuous: Perceptions of tissue damage are not pain unless they are represented *in* a pain experience. What's a pain experience? A representation of tissue damage.[75] The second possibility is also ultimately unsuccessful. It is clear that Tye thinks that tissue damage is *often* represented as bad, but he ultimately seems to say that badness is only a

[75]Tye is well aware that more is required.

contingent, typical feature of pain experiences.[76] Third, Tye takes all pains to have some "common phenomenological core."[77] I take this to be the most plausible option. This common core, however, receives only vague descriptions: the "characteristic sensory component," "a quality I strongly dislike," and "a certain sort of sensory-affective representational content." If *painfulness* is the intended phenomenological core necessary and sufficient for a representation of tissue damage to be a *pain* experience, then all of the arguments from the previous section apply directly. Consistent with his representationalism, moreover, the common core must represent *something*, but we are not told what.

Lacking clarity about Tye's characterization of what a pain experience *is*, it is difficult to assess how many of the dissociative cases canvassed above will be problems for his. The discussed cases of perception of tissue damage without pain, may fail to meet Tye's criteria for a pain experience. In episodic analgesia, for instance, perhaps there is no pain in a wounded soldier's leg because he does not represent *badness* and/or whatever objective quality corresponds to pains' "common phenomenological core." Episodic analgesia and its ilk may not then be problems for Tye. His proposal does not provide clear enough criteria to say.

Even if further clarification allowed Tye to account for perceived tissue damage without pain, his view cannot accommodate pain without the perception of tissue damage. These are all hallucinations or illusions for Tye, just as they were for Armstrong:

---

[76] E.g., Tye, ""Another Look at Representationalism about Pain," 106-107. It is also not at all clear what we are held to be representing when we *do* represent the "badness" of pain. Tye mentions the release of prostaglandins—an organic acid often released by damaged tissue. There is prostaglandin release in, for instance, pain asymbolia and episodic analgesia where there is no experience of awfulness. Perhaps in these cases there is just no *representation* of the prostaglandin release?

[77] E.g., Ibid., 99.

> Pains, viewed as experiences, intuitively can be misleading or inaccurate. Take the case of phantom limb pains. You have no right leg and yet you experience pain in the leg. How is that possible. Answer: in undergoing the pain experience, you *hallucinate* a right leg.
>
> Referred pains are also nonveridical. You can feel a pain in the left arm, when there is nothing wrong with the arm, the cause of the experience being a disturbance in the heart. Such a pain intuitively is inaccurate or misleading; for without additional information, on the basis of the pain, you would be disposed to nurse the arm, to rub it, to believe that something is awry in the arm itself. The case is one of *illusion*. (100)

As we have seen above, however, much more than phantom limb and referred pain are classified as hallucinations and illusions by Tye's perceptual view. Chronic pain and headaches, along with all the others canvassed, are deemed hallucinatory.

Tye's view is also vulnerable to additional dissociative cases that presented no problem for Armstrong, because he characterizes the affective aspects of pain as components of the perception of tissue damage. The same type of tissue damage (typed by cause and severity) undergone by different people and the same person at different times, however, is known to involve different affect and different emotions.[78] Consider the difference in experienced affect of a person who habitually engages in self-injurious behavior when cutting their arm, and the yet different affect if they cut their arm unintentionally. If you drill a hole in my head and scrape my skull, I will not sit calmly and catch the falling blood as the patients of Trepanation routinely do. The examples need not be so dramatic and they are discussed in every pain textbook. I may skin my knee without it feeling all that bad, while the pain from your skinned knee is excruciating.

About all these cases, Tye thinks that, frankly, some people are just *better* at representing tissue damage. This should not surprise us, he thinks, since we face the same situation with the other modalities. He argues that our brains were designed by evolution to, among other things,

[78] See, for instance, McMahon et al. (2006). That this is common currency would, however, be supported by any contemporary pain textbook. Moreover, I take it that it is common knowledge.

represent objective qualities. Evolutionary design provides the optimality conditions for those representations—though we may never know them. What we *do* know, according to Tye, is that when the same objective quality is differently represented, someone isn't functioning optimally.[79] Every type of tissue damage has an optimal representation determined by evolution—variation from which is not veridical. This includes whatever objective quality of the tissue damage the affective aspect purportedly represents in a pain experience. According to Tye, if knee-skinning hurts you more than me, one of us is undergoing an illusion. This will, quite obviously, make many, *many* pain experiences illusory.

Perhaps the immensity of hallucinations, illusions, and misperceptions is just something we should accept. On Armstrong's view: Tactition and bodily perception are subject to *many* illusions and hallucinations and often fail to transmit information about the sensible qualities distinctive to those modalities. On Tye's view: We undergo vastly more illusions and hallucinations about pain than we thought or can ever be sure of. Query Armstrong: Why are some sense-impressions as of bodily disturbances at locations pain when others are not? Query Tye: Why do some perceptions of the same type of tissue damage at locations yield excruciating pain while others yield none at all? Answers: No modality is perfect.

This is implausible. A characterization of pain, like Armstrong's, as a sense-impression abstracted from immediate perceptions of tissue damage deems too many states illusions and hallucinations, and characterizes too many states as pain that are *not* pain. Armstrong himself was open to the idea that his characterization of pain may require revision in response to these pressures. A characterization of pain as a perception that represents tissue damage and its qualities in an ill-defined "pain experience," like Tye's, deems an abundance of pains

[79] As applied to pain in particular see, for instance, Tye, *Reply to Commentaries,* 171.

hallucinations and illusions. If badness is the "common phenomenological core" whose representation makes the experience a *pain* experience, then *most* pain experiences are hallucinations and illusions. Though Tye does not appear concerned about this, he should be. While there are no strict rules about how many illusions, hallucinations and failures of perception are too many, I think the evidence is compelling that if we characterize pain as a perception, Tye's view, like Armstrong's before him, will saddle us with too many. The second most common type of illness (chronic pain) would be a hallucination.

What *is* plausible about perceptual views is that pain often *includes* a perception (as) of tissue damage. Why, however, should we try to shoehorn all pains into requiring this perception *or* as decreeing it sufficient? If we characterize pain as having multiple components, no one of which is necessary or sufficient, the plethora of dissociations between pain and bodily disturbances are no longer problematic. At the same time, a complex account of pain preserves the fact that many pains *include* a perception of, or (more cautiously) a sense-impression as of, tissue damage at a location. No one should deny this. Cases lacking any perception of tissue damage at locations are still pain, however, if the other paradigmatic components of pain are present. Cases including perception of tissue damage at locations are *not* pain if the other components of paradigmatic pains are *not* present.

Consider how well a complex account preserves our intuitions about who is and is not in pain validates patient reports. The chronic pain patient is in a state with negative affect, with a specific, qualitative content typical of pain, and exhibits characteristic withdrawal and avoidance responses. As such, the state is pain. The soldier on the battlefield, though perceiving bodily damage, is *not* in a state with negative affect, does not report experiencing qualitative content

typical of pains, and does not exhibit characteristic withdrawal and avoidance responses. As such, the state is not pain.

It is worth considering whether a unitary perceptual view could avoid these is characterized as a distinct modality. Following MacPherson (2011), I take perceptual modalities to be individuated by four main criteria: proximal stimulus, representation, phenomenal character, and a dedicated sense organ.[80]

Using these criteria, pain is unlikely to be a distinct modality. First, pain has no distinctive proximal stimuli. The best candidates are thermal, chemical and mechanical stimuli, but these are also the proximal stimuli of tactition. Second, nociceptors are probably the best candidate for a distinctive sense organ, but this suggestion faces two well-known problems. First, on mainstream views, wide-dynamic-range cells (WDRs) are necessary and sufficient for transmitting nociceptive information. Necessity is advanced on the grounds that they receive information in the dorsal horn that must be passed along to the brain. Sufficiency is advanced on the grounds that WDRs carry nociceptive information after nociceptor activity has been shut down (e.g., in cordotomy). WDRs are, however, non-controversially the transmitters for tactition—they are not a distinctive sense organ for pain.[81] Second, 'nociceptor' is largely heuristic (Craig, 2003b). No type of afferent is unresponsive to innocuous stimulation, and those preferentially responsive to innocuous stimulation are often responsive to noxious stimulation. This renders nociceptors a loose subset of receptors that, in general, respond to all manner of thermal, mechanical, and chemical stimulation subserving, again, tactition and not just pain.

---

[80] As MacPherson acknowledges, these criteria will be more or less independent depending on one's adopted theory of each.

[81] See Melzack and Wall's (2008) gate control theory and Craig's (2003a, 2003b) homeostatic emotion theory for more on the relationship between WDRs and nociceptors. Chapter 3 of my dissertation discusses these in some detail.

Concerning what is represented, we have seen above at length that it is very hard to find a well-correlated objective quality; tissue damage is the best candidate, but the correlations are tenuous.

Whether the phenomenal character of pain can be employed to establish pain as a distinct modality, however, deserves further discussion. If we weigh this single criterion of modality individuation heavily enough, then we may choose to sidestep the problems with the other three criteria. I argued in the previous section that pains are not well characterized by any distinctive quale, but left open the possibility that a quality-space theory of pain may do a better job. If we could determine a pain quality space (as against simply a quality space for tactition or bodily perception), then we may want to conclude that pain is a perception after all.

The dimensions of any quality space of interest are determined, for a creature, by what that creature can discriminate. Austen Clark, in "Painfulness is Not a Quale," walks through a sketch of the steps for constructing a pain quality space. He notes that we first "… need to canvas the content of all the similarities and differences one can possibly sense among episodes of pain."[82] The goal is to identify the dimensions of the quality space such that *every* discriminably distinct pain has a *unique* location within the space. We identify the individual dimensions of the space by identifying which qualities of the creature's sensations vary independently of the others. Each dimension, furthermore, should correspond to some class of perceptible properties.[83] Ideally, the similarities and differences of the perceptible properties perceived through the modality of interest should be isomorphic to the similarities and differences of the relevant mental qualities that enable the perceptual discriminations.

In sum, constructing a quality space for pain requires: (1) Determining the dimensions of pain's quality space, such that each discriminably different pain has a unique location in the

[82] Clark, "Painfulness is Not a Quale," 182.
[83] Cf. Clark (1993) and Rosenthal (2005).

space; and, (2) determining the perceptible properties that are appropriately correlated with the identified dimensions. Both steps in constructing the space are problematic.

First, beyond felt location and intensity, there are no good candidates for the dimensions of the quality space. The McGill Questionnaire is, to date, our best way of characterizing the contents, including the qualitative contents, of pain. Nonetheless, it is poorly suited for providing dimensions of a quality space; a single pain experience is often described by multiple descriptors from the same group. If we use the questionnaire, discriminably different pains will be uniquely located in the space.

Second, and perhaps more problematic, what perceptible properties are these elusive dimensions isomorphic to? There are two good candidates: the qualities of tissue damage and the more general properties of the primary afferents. Despite its intuitive plausibility, however, tissue damage and its qualities are not well correlated with pain experiences and their qualities. This is, I have argued, the downfall of the other perceptual accounts. In light of this, the best candidates are the perceptible properties of the primary afferents: temperature, mechanical pressure and chemical stimulants. If we go this route, however, then we will simply end up with the same quality space for pain as for tactition. A view that characterizes the qualitative content of pains by a quality space for tactition is obviously not a view that characterizes pain as a distinct modality.

It may seem that my objections to a quality-space theory for pain amount simply to throwing up my hands: I do not know how it can be done, so it must not be doable! The fact is, however, that I *don't* know how it can be done, and neither does anyone else. All of the suggestions in the literature take us back to tactition or dissolve into head scratching. Such an objection is admittedly not very powerful by itself.

Consider also, however, that a quality space for pain was supposed to provide strong enough reason to dismiss the problems raised by the other three criteria for modality individuation. Without a well-founded quality space for pain, this goal remains unmet. We are unable to fulfill any of the four criteria

There is no clear way, and more important no *reason* to develop a quality space for pain. This is not to say that quality space theories are not a good way of characterizing qualitative content (or better here: character). Indeed, I think they are the best way.[84] The best candidates for the perceptible properties corresponding to the dimensions of a pain quality space are the same perceptible properties that characterize the qualitative content of tactition and bodily perception (temperature, chemical stimulants, and mechanical pressure). I believe that we ought to work on constructing a quality space for tactition and, perhaps, for bodily perception.[85] Insofar as those quality spaces do not provide adequate characterizations of the differences between each discriminable pain, however, the moral is that there is more to pain than perception—not that pain is a distinct modality.

In sum: Pains often include perception, but no perception is necessary or sufficient for pain. The best candidate for what is perceived is tissue damage and qualities of that damage. The cases canvassed above show that correlation between the perception of tissue damage and pain are poor; too many non-pains are classified as pains, and too many pains are classified as hallucinations and illusions. Perhaps we would just have to accept this if there was no other theoretical option, but there is: the perception of tissue damage is just one paradigmatic

---

[84] I will not argue for a quality space theory here. My point is that pains are not best characterized as a perception even if we accept a quality space theory of mental qualities.
[85] It is not clear to me whether we should think of these as distinct modalities. Armstrong, as above, thought that tactition always also involved bodily perception. None of my arguments turn on this.

component of the complex experience of pain. We should accept the clinical evidence and admit that there are many pains without any such perception, and that there are many cases of the perception of tissue damage without pain. These cases undermine the claims that perception of tissue damage is necessary or sufficient for pain, and the (weaker) claim that pain is adequately *characterized* simply as a perception. We may try to distance the perceptual view from the dissociative cases by classifying pain as a distinct modality, but such a view is entirely lacking in motivation; none of the criteria for modality-individuation are met. Pain is not simply a perception.

4. Pain is Not Simply an Emotion

Despite disagreement about what emotions *are*, paradigmatic pains involve an aspect that most philosophers consider an emotion. Unitary sensation and perceptual views do not recognize this emotional aspect to be a *component* of pain when present. This is mistaken; pain experiences paradigmatically include an emotional component. A mistake in the opposite direction would be to take the emotional aspect of pain to be not only constitutive when present, but exhaustive or essential to pain. There are no known unitary views of pain as an emption.

Token pains, at least *prima facie*, vary in emotion or affect. In addition to the ten groups of sensory descriptors provided in section 2, The McGill Pain Questionnaire provides five groups of descriptors for the affective component of pain (11-15) and an evaluative group (16) that is, arguably, also emotional:

Group 11: Tiring, Exhausting
Group 12: Sickening, Suffocating
Group 13: Fearful, Frightening, Terrifying
Group 14: Punishing, Grueling, Cruel, Vicious, Killing
Group 15: Wretched, Blinding
Group 16: Annoying, Troublesome, Miserable, Intense, Unbearable

As with unitary sensation views, a unitary emotional view privileging any single one or combination of the emotions described above will be problematic. Some token pains will lack the identified emotion(s), and the emotional component of a token pain often varies over time.

A more general emotion that all token pains share in common may be proffered instead. As with the unitary sensation views, the best candidate appears to be painfulness—now understood as an emotion. As concluded in section 2, however, there is robust clinical evidence that some pains are not painful, and this remains true whether painfulness is characterized as a quale or as an emotion. No other candidates for a distinctive pain-emotion suggest themselves.

The emotional dimension of pain does receive special attention by Clark (2005), Gustafson (2005), and Craig (2003a, 2003b, 2003c). A quick survey of these three views reveals that all three theorists ultimately reject unitary accounts of pain in favor of more complex accounts. Emotions are emphasized as, but *only* as, an important component of pain.

First, consider Gustafson (2005). In his article, "Categorizing Pain," he makes two central claims: (1) Pain is not a simple sensation; and, (2) Pains are best characterized as "emotionlike," "needful," or "affectivelike" states. Gustafson argues that we should categorize pain however will best serve scientific research and clinical practice, and this has, and is expected to be, revised over time. Construing pain as a simple sensation is not only implausible to Gustafson, but he argues that it is scientifically impotent. Characterizing pain as an emotion-like state, on the contrary, is argued to be our best bet for theoretical and practical dividends as evidenced by contemporary pain research.

Gustafson makes plain, however, that he does not think this emotion-like state is properly characterized as any *one* affect. He considers a number of possibilities for a general, essential, emotional, component of pain and concludes that they all fail. Of the possibility, for instance,

that pains should be categorized as those mental states that are *unwanted*, he says: "If it is claimed that the sole pain quality is the feature of being unwanted by the host, we should remember that not all injuries associated with pain are unwanted."[86] He similarly observes that characterizing pain as being a *hurtful*[87] state is equally overly simplistic: "… hurtfulness, too, does not appear to be a single feature … If we reduce it to a single parameter, as we might in theory, it will pretty clearly cover more than pain."[88] There is no single affect that Gustafson accepts as an adequate characterization of pain.

Gustafson, then, appropriately eschews any unique emotion as characterizing pain, but he then errs in seeming to take some affective component *or other* to be necessary. Quoting the offending passage at some length:

> Clearly the pains experienced are not simple sensations arising from free nerve endings in the skin. Whatever the full story about them, it is not a simple case of transduction, transmission, and qualitative sensation. *Affective components are not separable from the pain experience*. Cognitive set or schema can affect such pain experiences. The (statistical) intercorrelations of these aspects… make it impossible to think of them as separable parts of a whole. Only under drug treatments of various sorts, or brain damages, and possibly hypnotic conditions can they dissociate, and then only relatively. They appear in experience as emotional distress (234, emphasis added).
>
> And again:
>
> [T]hese data also suggest that the category most likely to be useful is among the emotion, motivation, and need categories. In fact, pain researchers in the pain sciences have for some time urged this new classification. Philosophers have lagged behind the sciences in this regard" (234).

Insofar as Gustafson takes pains to be *necessarily* emotional or affective, I think he is in error. I am surprised that Gustafson sometimes sounds as if he takes this position given his recognition of the many dissociative cases presented in the pain literature and his stated desire to harmonize

---

[86] Gustafson, "Categorizing Pain," 224.
[87] 'Hurtful' here differing only negligibly from my 'painfulness.'
[88] Ibid., 224.

with contemporary pain science. The affective or emotional component of pain is there presented not only as occasionally dissociating, but also as *often* dissociating. Mild pains, pains undergone when distracted or in danger, pains resulting from self-inflicted damage, the pains of the masochist, the pains of the mystic and ascetic, and more, are all cases of the sensory aspect of pain uncoupling from its affective aspect. These extend beyond the brain damaged, drugged and hypnotized.

Despite this mistake and more relevant to this chapter, Gustafson rejects unitary accounts of pain. A substantial part of Gustafson's article consists in his approvingly quoting the complex views of pain that I take up directly in the next chapter. After quoting at some length, he notes: "I could continue to cite sources from recent pain research … showing the move away from pain as sensation toward categorizing pain as a complex, emotional and motivational condition."[89] Earlier, he asserts bluntly that "… it is best to regard pains as complex emotional and perceptual processes or events."[90] Though Gustafson recommends thinking of pain as an emotionlike state, his view should not be confused with a unitary account that construes pain as being *simply* or *only* an emotion.

Similar comments apply to the account presented by Austen Clark (2005). Clark characterizes pain as a "primitive feeling state"[91] and assimilates it to other such identified states including hunger and thirst. Though he calls these states *feeling* states, Clark's account is decidedly not a unitary account. All primitive feeling states are to be understood on a "tandem model" that posits two necessary components: a sensory component and a motivational component. For pain in particular, the relevant sensory component is identified as one of "some

[89] Ibid., 237.
[90] Ibid., 233.
[91] Clark, "Painfulness is Not a Quale," 191.

range of qualities within those discriminable by nociceptive somesthetic discrimination,"[92] and the relevant motivational component is painfulness.

Clark construes painfulness broadly. He rightly notes that pains do not share any *particular* motivation or emotion in common. To be painful, a state need only have *some* motivational element to avoid, and this motivational element need only be dispositional: motivations are aroused by the sensory component of the feeling states, but need not be acted on and may be trumped by other, stronger motivations[93] Summing up his motivational account of what makes a state painful:

> … the motivational change might show up … as a change in the content of desires, preferences, inclinations, incentives, reinforcers, and so on. The aversion might show up in any of these ways; no single one of them is essential. What is essential to painfulness is that it arouses one or another of them. Another way to put this: if *S* is painful, then *S* arouses some disposition to avoid. Avoidance, like aversion, is a catch-all term. It indicates the general direction of one's motivations… (186).

Painfulness is some motivation or other to avoid.

Primitive feeling states are just those states where sensations and motivations are biologically linked. Such states are not simple sensations, but they are also not simple desires or emotions. They are, however, things that we *feel*. According to Clark, we *feel* pain when we have a sensory experience characterized in the range of nociceptive, somesthetic qualities, *and* have some aversive, motivational disposition about what is sensed.

Despite their biological linkage, Clark takes the identified sensory and motivational components of pain to be logically distinct, and to dissociate in actual cases.[94] About whether these dissociative cases are *pain*, however, he is somewhat equivocal: "It feels just like a pain, but it is not awful; it does not bother you; its continuation or cessation is a matter of indifference.

---

[92] Ibid., 189.
[93] Ibid., 186.
[94] Ibid., 177.

Would that state be pain? Well, Socrates, yes and no."[95] "Yes" is because the sensory component is present, and "no" because the motivational component is absent.

Despite recognizing the dissociative cases, Clark ultimately seems to conclude that our ordinary concept of pain *appropriately* applies only when both of his identified components are present. As he says plainly: "The ordinary term 'pain' refers to a state of affairs constituted by both [sensation and motivational responses], standing in a certain relation to one another."[96] Similar claims are made a number of times toward the end of the article.

Clark's account is laudable insofar as it rejects unitary accounts of pain, but errs insofar (which may not be far) as it requires any one component to be necessary to pain. His essentiality claims are careful and subtle, and perhaps he does not mean to go so far. Regardless, for purposes here, it is enough to conclude that though Clark calls pain a primitive feeling state, his view should not be confused with a unitary account that construes pain as being *simply* or *only* an emotion.

We turn finally to Craig whose view can helpfully be understood as the neuroscientific filling-in of Clark's philosophical outline. The states that Craig calls *homeostatic emotions* are similarly characterized and appear to be coextensive with Clark's primitive feeling states. Craig's inspired, new theory will be taken up in some detail in the next chapter, which focuses on complex accounts. For now, I will just make clear that Craig's account of pain is anything but unitary. Consider the following passage:

> Pain is an enigmatic feeling from the body, distinct from the classical senses because it is multifaceted (it is a discriminative sensation, an affective motivation, a potent autonomic drive and a reflexive motor stimulus) and because it is inherently variable ("A New View of Pain as a Homeostatic Emotion," 303).

---

[95] Ibid., 188.
[96] Ibid., 190.

Craig clearly considers pain to have multiple components.

Though Craig characterizes pain as *an* emotion, this emotion is not a single mental state, but multiple states bearing complex relations to one another. On Craig's view, emotions just *are* complex processes involving multiple components and sub-served by multiple mechanisms. This leads Craig to conclude that "[e]motions consist of a sensation and a motivation with direct autonomic effects, and in this new view, pain is one of many distinct homeostatic emotions that directly reflect the condition of the body."[97] Again, the details of Craig's view will be taken up in the following chapter. For now, we can conclude that though Craig calls pain a homeostatic emotion, his view should not be confused with a unitary account that characterizes pain as being *simply* or *only* an emotion.

The central thesis of *this* chapter is that no unitary account of pain adequately characterizes the phenomenon. Neither Gustafson's, Clark's, or Craig's emotional accounts of pain—even if correct—present a challenge to that thesis.

5. Conclusion

In chapter one, I made the following promissory note: "Paradigmatic pain experiences, I will argue, include several aspects, and isolating any one of these aspects is an inappropriate characterization of the phenomenon. Views that isolate some one aspect commit the error of under-specification." We have now considered the unitary accounts of pain and seen that they do, indeed, commit the taxonomic error of under-specification.

The demise of all of the unitary accounts considered here, and of any that may be offered in the future, is the same: the wide range of pain phenomena are not adequately characterized by

[97] Craig, "A new view of pain as a homeostatic emotion," 304.

any privileged one of its aspects. A paradigmatic pain includes sensation, perception, emotion, cognition (thoughts), and motivational responses. These multiple dimensions of pain are only contingently connected so it should not be surprising that they dissociate, though it may be surprising that they dissociate so frequently. As we have seen and will continue to explore in the following chapters, numerous pain experiences lack each of these, and numerous experiences include any one of them without being pain. Whichever dimension a unitary account privileges, an unacceptably large number of pain experiences are left out and an unacceptably large number of non-pains are included. Paradigmatic pains involve multiple components and no one component is necessary or sufficient.

The next chapter focuses on composite accounts. Though no philosophical, unitary account of pain adequately captures the wide range of pain phenomena, perhaps these complex, scientific accounts can. Recall from chapter 1 that to establish that pain (or anything else) is a natural kind, requires not only an adequate characterization of the phenomenon, but an identification of some underlying, homeostatic mechanism that explains those properties by which we characterized it. In considering the unitary accounts, there was no need to go on to evaluate whether or not the identified phenomenon was explained by a single underlying mechanism; shoehorning all pain experiences into one type of mental state did not adequately characterize the phenomenon to begin with. Complex accounts characterize pain more adequately by recognizing its multiple components. I argue in the following two chapters, however, that the resultant complex is nonetheless not explained by the workings of a single, underlying mechanism. There are mechanisms for sensations, perceptions, emotions, and motivational systems, but there is no single mechanism that underlies the combination of these that constitutes any token of pain.

Chapter 3: The Inadequacy of Complex Accounts

1. Introduction

Let's review. In chapter 1, I developed an account of natural kinds as those kinds usefully referred to in the generalizations of the relevant science, and argued that they will submit to a mechanistic explanation in principle; i.e., a kind we can show of such will *not* submit to a mechanistic explanation is not natural. I claimed that chapters 2 through 4 would establish that pain would not submit to a mechanistic explanation. In Chapter 2, I argued that all unitary accounts under-specify pain and, as a result, fail to do justice to the wide range of pain phenomena. In this chapter and the next I argue that all complex accounts of pain, including those currently on offer by pain science, fail to yield a natural kind because they lump together many phenomena not explained by a single, underlying mechanism. Recall that the claim that pain is not a natural kind is appropriately broken down into two: 1) Pain is not a natural kind; and, 2) No type of pain is itself a natural kind. Accordingly, the current chapter focuses on complex accounts of pain as such, and the next chapter focuses on complex accounts of types of pain.

I consider five models of pain currently dominating pain research. In section 2, I provide a somewhat detailed overview of Melzack and Wall's colossal gate-control theory, followed by briefer overviews of Price's innovative somatic perception theory (section 3), Craig's more recent homeostatic emotion theory (section 4), the elaborate pain neuromatrix theory (section 5), and, finally, the definition of pain offered by the International Association for the Study of Pain (section 6). While these complex accounts are better able to accommodate the wide range of pain phenomena, they do not identify a pain mechanism that can explain the many types of pain. Some of these theorists are explicit that they have *not* done so.

The failure of the five dominant models of pain to yield a natural kind presented in this chapter may convince some of pain's non-naturalness, but others may maintain that a superior model of pain that identifies a mechanism explaining it should still be hoped for and sought. This worry will be fully diffused only in the following chapter: inquiry into different types of pain will show *why* the wide range of pain phenomenon will not submit to a single mechanistic explanation. Different types of pain implicate different mechanisms, and we now know that some mechanisms are involved in some pain experiences and not others. An explanation of all pains by a single, underlying mechanism is not forthcoming. Claim 1 will thus be established; pain is not a natural kind.

2. Melzack and Wall's Gate-Control Theory

Appreciating the significance of Melzack and Wall's gate-control theory requires a few words on the history of pain science. One approach to that history is through the debates among the scientists, and two of these are particularly relevant.

First is the traditional debate about what *sort* of thing pain is. A common way of presenting the debate within science is as of a competition between three main theories: the intensity theory, the sensory theory, and the affective theory. According to the intensity theory, pain is best understood as a perceptual experience, but one that arises through *any* modality if the input is strong enough. By contrast, the sensory theory characterizes pain as a distinct perceptual modality, subserved by a specialized system. The affective theory of pain classifies pain as a particular affect or quale; it is perhaps best understood as construing pain as the extreme end of a spectrum whose opposite end is pleasure. Competition among these three theories was fierce and is as old as Aristotle. As Dallenbach (1939) describes this debate in his much respected article:

> The intensive theory and the sensory theory stood in opposition to one another, and they both stood in opposition to the traditional pleasure-pain theory which represented pain as an affective *quale*. In the decade between 1886-1896, these different opinions clashed and a three-cornered controversy ensued, the like of which has never before, nor since, appeared in the scientific literature (337).

By the time he writes his history, Dallenbach takes the intensity theory to have been completely disproved, the affective theory to have been recast so as not to be a contender, and the sensory theory to be partly established and well on its way to completion. He concluded, accordingly, that pain is a distinct modality. The evidence seemed to him incontrovertible that pain was subserved by dedicated, specialized fibers and that dedicated, specialized pain pathways were being discovered even as he wrote. It was just a matter of time, Dallenbach and most working scientists of the day thought, until the specialized pain mechanisms were fully uncovered.

The second, though related, line of debate concerns whether pain should be understood as a specific phenomenon or a convergence of many. Craig's (2003b) presentation of the history of pain science, for instance, focuses on this debate. The issue is whether pain is a "distinct sensation represented by specialized elements both peripherally and centrally," or rather an "integrated, plastic state represented by a pattern of convergent somatosensory activity within a distributed network."[98] Species of the former view are called specificity theories, and species of the latter, convergent theories. Notice that these two conceptions of pain crosscut with the three theories presented in Dallenbach. Whether pain is a perception or an affect, is it subserved by specialized mechanisms or by the convergent activity of multiple mechanisms no one of which explains it?

[98] Craig, "Pain mechanisms: labeled lines versus convergence in central processing," 1.

Before Melzack and Wall, the leading theories were all specifity and sensory theories. Is there a pain center in the brain? Are there receptors just for pain? Are there pathways specialized for pain processing? Is pain a distinct modality? The leading answers were affirmative.

Melzack and Wall's resounding negative answers to *all* of these questions permanently changed the scientific study of pain. The biggest insight and influence of their gate-control theory is the claim that the phenomenon that we call pain results from the activity of multiple, converging mechanisms, no one of which is specific to pain. The theory revolutionized pain science:

> The revolution began with a radical shift away from the traditional concept of a specific pain pathway to a brain centre where pain sensation is felt—a single sensory dimension that varies only in intensity. Instead, the gate-control theory postulates that pain is a unified stream of experience that is generated by the brain and is influenced by all of its cognitive functions such as memories of prior experiences and the meaning of the current situation ("*The Challenge of Pain*," X).

The spark that starts the fire is the denial that pain is a distinct modality; it is not subserved by specialized receptors, specialized pathways or specialized cortical areas. Melzack and Wall straightforwardly deny that there is any such thing as a 'pain centre': "… the old concept of a 'pain centre' is obviously nonsense. Many areas of the brain are involved in pain processes and they interact extensively."[99] The gate-control theory instead characterizes pain as an experience resulting from the wildly complex interaction of multiple systems. In sum:

> It is now becoming increasingly evident that virtually all of the brain plays a role in pain. Even seemingly unrelated brain activities such as seeing, hearing and thinking are important. Seeing the source of injury, hearing the sounds that accompany a rifle shot or a falling beam, and thinking about the consequences of an injury all contribute to the pain. Any satisfactory understanding of pain must include all of these processes which interact with inputs from the injured area or from deafferented neurons that produce pain signals when injury is absent (*The Challenge of Pain,* 132).

[99] Melzack and Wall, *The Challenge of Pain*, 122.

This convergent model now dominates the field and has legitimized research into the effects of almost *anything* on pain—fostering new, previously unimaginable, treatments.

Melzack and Wall's theoretical starting point is the same mismatch between tissue damage and pain that proved fatal to the perceptual theories discussed in Chapter 2:

> This is the essence of the puzzle of pain. Why are pain and injury not always related? What activities of the nervous system intervene between injury and pain perception that make the relationship so variable? (*The Challenge of Pain,* 3).

Their theory was explicitly devised to explain not only pain caused by injury or tissue damage, but the many kinds of pains *not* involving tissue damage. Further, they explicitly claimed that an adequate explanation of pain should explain not only why we sometimes feel pain without tissue damage, but also why we sometimes undergo tissue damage without feeling pain. Ultimately, the lack of correlation between pain experiences and *any* distinct activity in the body led Melzack and Wall to claim that specifity theory, in any of its forms, was untenable. They reasoned persuasively, and influentially, that if pain were the result of specific receptors, pathways, and cortical centers, then surgical intervention or lesions of those receptors/pathways/centers should abolish the pain. It did not.

Rather than discard the ideas of their predecessors entirely, however, Melzack and Wall accepted that the mechanisms posited by specificity and sensory theories were indeed implicated in pain experiences:

> However, when all the theories—from specificity theory onward—are examined together … it is apparent that each successive theory makes an important contribution. Each provides an additional mechanism to explain some of the complex clinical syndromes or experimental data that were previously inexplicable (*The Challenge of Pain,* 164).

Previous theories should not be rejected—they should be supplemented. While inadequate to explain the wide range of pain phenomena, the theories are not strictly incorrect, but incomplete.

Gate-control theory, too, is presented as *a* mechanism implicated by pain, and its postulation is needed to explain pain experiences that no previous theories could explain.

In the final presentation of the view (2008), gate-control theory is presented as a five-stage pain mechanism. Below is a brief overview of each stage, including developments of the theory between its original and final presentations.

- Stage 1

Stage 1 is basically specificity theory. Two types of small-diameter nerve fibres, Aδ (myelinated) and C (unmyelinated), deliver impulses to the spinal cord when stimulated. These small-diameter fibres have high-thresholds (they preferentially respond to relatively intense stimulation), making them the most likely candidates for pain-specific receptors, or "nociceptors". Their input is received in the dorsal horn of the spinal cord by "transmission" cells, often simply called 'T cells.' These T cells, in turn, transmit the received information to both local reflex circuits and on to the brain. As can be seen in his classic depiction of the pain mechanism, "[t]his stage incorporates all that is shown in Descartes' concept of pain."[100]

[100] Ibid., 166.

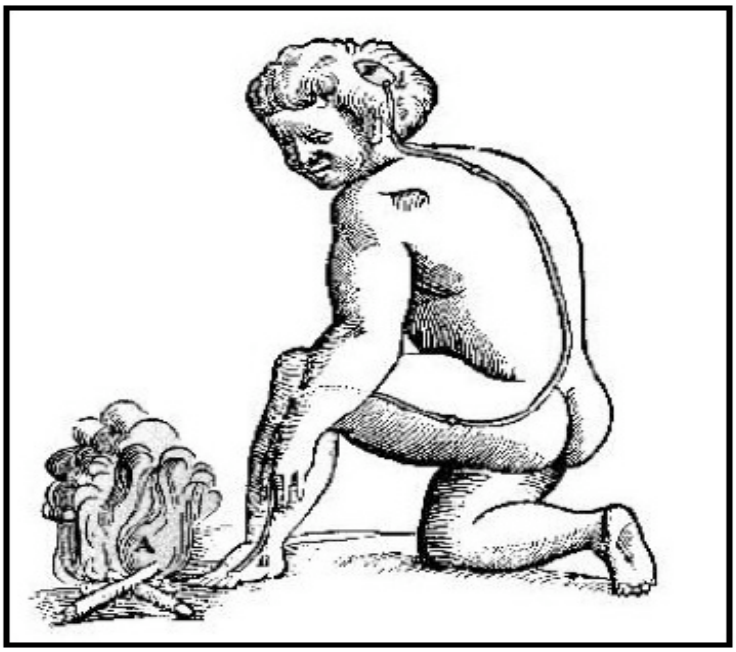

Figure 3.1:
Descartes' Pain Model

All that is added to this famous version of specificity theory are the T cells in the spinal cord.

Since gate-control theory's inception, our knowledge has advanced concerning not only the functionality of both small-diameter fibres (included), but large-diameter fibers (not mentioned above). Of particular relevance is that the functionality of the fibers is much more complex than previously thought, and the correlation between their stimulation and a pain experience is even more variable. By 2008, Melzack and Wall no longer thought that even the small-diameter fibres ("nociceptors") were usefully grouped together, claiming instead that "… new data… show that myelinated [Aδ] and unmyelinated [C] fibres have totally separate functions…"[101] A key difference is that C fibres become desensitized to repeated stimulation (it takes more and more intense stimulation to activate them), whereas Aδ fibres become more sensitized (it takes less and less). This difference implicates different roles for C and Aδ fibres in distinct types of pain experiences (e.g., in chronic pain). Another relevant discovered difference is that these fibres do not terminate in the same place in the dorsal horn. The receiving T cells

[101] Ibid., 171.

were originally posited as being in the substantial gelatinosa (lamina I and II), but the transmission cells themselves have been somewhat discredited as a unified class. Melzack and Wall eventually conclude that "… there is not a hint of a simple, straight-through transmitting cell."[102] As discussed below, the functionality of the T cells in the gate-control mechanism makes mismatches between pain and Aδ and C fibre firing to be expected. Increased understanding of the primary afferents (Aδ, C and the large-diameter Aβ fibres) only further confirmed these mismatches.

Melzack and Wall deny that specificity theory can account for the now well-established mismatches:

> First, the onset of firing does not coincide with the onset of pain. Second, the time course of firing relates very poorly to the time course of pain. Third, identical pains evoked by different stimuli such as punctate pressure or heat are associated with strikingly different patterns of activity in afferent fibres. Specifity theorists carefully avoid discussing pathological pains, which, they admit, do not follow the rules of the normal nervous system. However, it seems totally unsatisfactory to explain 'normal' pain and leave the pains which matter, those of the suffering patients, for a later time and another explanation. Our conclusion from the experimental results as well as from patients' reports, is that pain is the result of an analysis by the central nervous system of the entire situation at the time, taking into account the firing of nociceptors as well as other data at its disposal. (*The Challenge of Pain,* 171.)

Their conclusion in 2008 was that variability between nociception and pain had become more established than before, putting another nail in the specificity coffin.

- Stage 2

Stage 2 posits excitatory interneurons in the dorsal horn that facilitate transmission from the small-diameter fibres to the T cells. The posit was needed to explain the protracted firing of the T cells after only brief transmission from the small-diameter fibres. It was later found to explain a related and robust phenomenon called wind-up: Short, brief transmission from the small-

[102] Ibid., 171.

diameter fibres at regular intervals results in increasingly exaggerated responses by the receiving T cells. This often results in pain that is y disproportional to the input from the small-diameter fibres. Note again variation between nociception and pain.

- Stage 3

Stage 3 posits that large-diameter, or Aβ, fibres, also transmit to the T cells. It was initially thought that T cells were nociceptive-specific (NS) cells; i.e., that they could receive input only from high-threshold (small-diameter) fibres. Further investigation revealed that this was not so. Instead, there are three kinds of T cells in the dorsal horn: NS cells, wide-dynamic range (WDR) cells, and an unnamed group that responds only to low-threshold input. These WDR cells can receive input from both high-threshold (small-diameter) and low-threshold (large-diameter) fibres. Melzack and Wall consider WDR cells to be "… by far the most common of the centrally projecting cells which signal injury."[103] It is worth emphasizing that WDRs are both crucial for the gate-control mechanism and *not* specific to pain processing. This marks a major departure from specificity theory:

> The specifity theorists concentrated on the NS cells as the only cells that could be involved in triggering pain. We, on the other hand, simply proposed that pain would be triggered if the firing rate of any group of cells exceeded a critical level determined by the properties of the brain. (*The Challenge of Pain,* 167.)

Note that whether or not the firing rate in the dorsal horn reaches the "critical level" for pain is not determined by local mechanisms in the dorsal horn alone. This foreshadows the massive, and hitherto underappreciated, contribution of the brain in pain processing.

Though there is now extensive confirmation of the dorsal horn's three transmission cell types, controversy remains about which to emphasize in explaining pain. Specificity theorists, of course, continue to emphasize the role of the NS cells. Settling this controversy would take us

[103]Ibid., 167.

too far afield here. Melzack and Wall cite a number of studies that they think support their claim that "…pain is related to the WDR cells rather than the NS cells…"[104] Important for our purposes is that WDRs are crucial for the gate-control theory, and it is the prospects of *that* theory for rendering pain a natural kind we are considering in this section.

- Stage 4

Stage 4 posits gate-control theory's most famous component: the (presynaptic) inhibitory cells acting on transmissions from large-diameter fibres. Most every synaptic region in the central nervous system contains excitatory (as posited in stage 2) and inhibitory interneurons that modulate transmission, and antecedent research had revealed that T cells were inhibited when large-diameter fibres were stimulated. Inhibitory cells in the dorsal horn receiving large-diameter impulses and modulating T cell transmissions were therefore posited. Notice that this means that large-diameter fibres can both excite and inhibit T cells—they have a "double effect." It is the inhibition vs. excitation of the T cells that determines whether the T cell will transmit to the brain possibly, in turn, resulting in a pain experience or not; i.e., inhibition vs. excitation determines whether the "gate" is open or closed. Thus the title 'gate-control.' It is inhibition, above, all that had such explanatory power and led to novel treatments.

In their 2008 presentation, Melzack and Wall present numerous developments in dorsal horn inhibition research. Highlights include:

- Low-level electrical stimulation that preferentially activates large-diameter receptors called TENS (for transcutaneous electrical nerve stimulation) is now used as a form of treatment. It is a method for "closing the gate." Consider rubbing your toe after stepping on a tack.

[104]Ibid., 173.

- Intense stimulation of *small-diameter* fibres is now thought to sometimes inhibit their otherwise excitatory effects on T cells via an inhibitory mechanism involving the brainstem called DNIC (diffuse noxious inhibitory control). In short, one intense stimulus inhibits another by way of the brain stem. DNIC may explain "counter-irritation"; e.g., why pinching yourself really hard can sometimes make your headache stop hurting.
- *Post*synaptic inhibition was initially only conjectured, and is now taken to be conclusively established. The relation between postsynaptic and presynaptic inhibition on the T cells, however, remains something of a mystery.
- T cells were discovered to also be inhibited by endogenous opioids, leading to the development of micro-injection of narcotics (morphine, in particular) into the spinal cord. The relevant mechanisms of action, however, are still mysterious.

Note that gate-control theory now has multiple available explanations for variability between nociception and pain that are unavailable to the specificity theorist.

- Stage 5

The fifth and final stage of the gate-control mechanism posits downstream influence on the inhibitory neurons in the dorsal horn from the brain, and a feedback loop from the spinal cord to the brain and back again to the T cell inhibitory neurons. That the brain exerted *some* influence was known well before gate-control theory. The loop was originally posited on the assumption that ascending information affects descending modulation, and despite *extensive* research, the brain's descending influence on the inhibitory neurons is still only poorly understood. The original model of the gate-control mechanism simply included an arrow from "central control" to the dorsal horn.

Though still mysterious, there have been some advances. It was discovered (Wall, 1967) that the T cells posited in the gate-control mechanism were subject to: "powerful, steady inhibition [that] flowed continually from the brain stem."[105] More and more descending, inhibitory, modulatory pathways from brain to spinal cord and back again have been discovered since. It has become clear that the idea of a *single* path of descending control on some *single* type of transmission cell is woefully over-simplified; there are many types of T-cells relevant to pain processing and *many* paths of descending control. As in the introductory quotes above, Melzack and Wall held that almost the entire brain exerts not only inhibitory, but excitatory, control over local and ascending dorsal horn transmission. All they think we currently have, however, are: "… very crude phenomena which can hardly explain the presence of an elaborate, interconnected, many-factored control mechanism,"[106] and much contemporary work continues focuses on descending modulation. The crude approximation of a "subtle two-stage mechanism" of descending control that Melzack and Wall presented in 2008 is:

> …1) the brain gives 'permission' to local circuits in the spinal cord to respond if the situation merits the responds, and 2) the local spinal cord circuits measure the input and trigger a reflex response if they are permitted to do so by the brain, and if the afferent input is sufficient (*The Challenge of Pain*, 176).

This model is a merely a sketch that needs considerable filling in.

In accordance with these five stages, the gate-control mechanism for pain presented in Melzack and Wall's original paper was represented as:

---

[105] Ibid., 169.
[106] Ibid., 176.

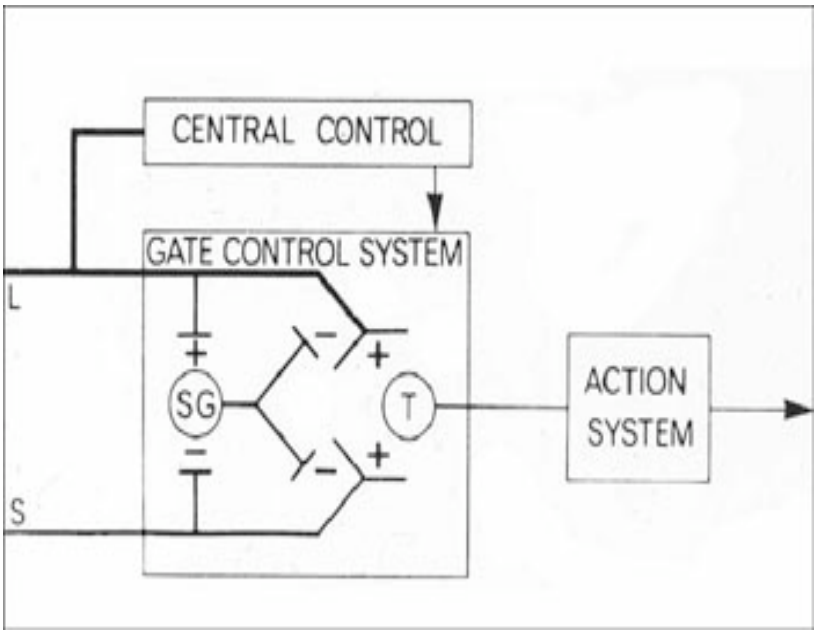

Figure 3.2:
Gate-Control Mechanism
Taken from: *The Challenge of Pain*

Where 'S' and 'L' represent large and small diameter fibres, 'SG' represents the substantia gelatinosa (lamina I and II), 'T' represents transmission cells and '+' and '-' represent excitation and inhibition, respectively.

Recall that Melzack and Wall were out to supplement the then-popular specificity theory because it could not account for a wide range of pain phenomena. The complexity involved in the gate-control theory, by contrast, made new explanations of a wide range of pain phenomenon available. Just as important, gate-control theory included multiple dimensions in its explanation of each token pain experience.

Melzack and Wall think it a major failing of specifity theory that it includes only pain's sensory dimension:

> The sensory approach to pain… valuable as it has been, fails to provide a complete picture of pain processes. The assumption that pain is a primary sensation has relegated motivational (and cognitive) processes to the role of 'reactions to pain'…. It is apparent, however, that sensory, motivational and cognitive processes occur in parallel, interacting systems at the same time. … [These] processes must be included in any satisfactory theory of pain. (*The Challenge of Pain*, 161.)

As intuitive as a purely sensory conception of pain might be, it is insufficient. Melzack and Wall lament that "[t]he concept of pain as purely a sensory experience long overshadowed the affective and cognitive dimensions of the total pain experience."[107] These additional dimensions are taken to receive not only phenomenological support (what it is reported as being like, subjectively, for a person in pain), but are predicted by the *multiple* ascending pathways from the spinal cord implicated in pain processing that were discovered throughout the 1950s.[108] These processes occur in parallel, interact with each other and other aspects of pain processing, and exert influence on whether the gate is "open" or "closed." Melzack and Wall conclude that it is illegitimate to characterize them as mere *reactions* to pain instead of genuine *components*. Specificity theory is then understood to be insufficient on three main grounds. It fails to explain the wide range of pain phenomena, it fails to explain the multiple dimensions of any token pain, and it cannot accommodate the multiple ascending pathways implicated in pain processing.

Pain is still discussed by Melzack and Wall as a perceptual experience—but it is a "perceptual experience" with many dimensions, only one of which is sensory. Their account is decidedly complex:

> It is assumed that these three categories of activity [sensory-discriminatory, affective-motivational, cognitive-evaluative] interact with one another to provide *perceptual information* regarding the location, magnitude, and spatiotemporal properties of the noxious stimulus, *motivational tendency* toward escape or attack, and *cognitive information* based on analysis of multimodal information, past experience, and probability of outcome of different response strategies. All three forms of activity could then influence motor mechanisms responsible for the complex patterns of overt responses that characterize pain. (*The Challenge of Pain,* 192.)

Accordingly, the final version othe gate-control theory is represented as:

---

[107] Ibid., 191.
[108] Ibid., 131-137.

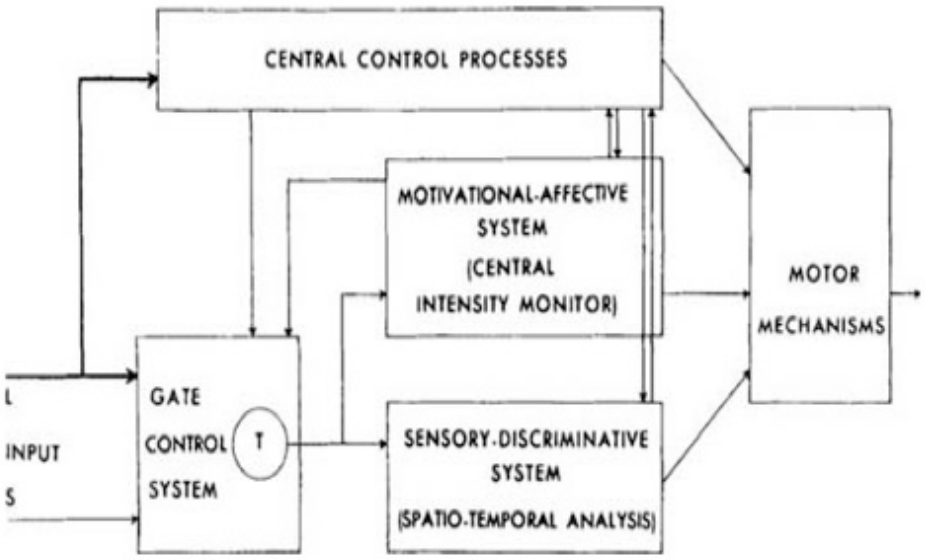

Figure 3.3: Gate-Control Theory. Taken from *The Challenge of Pain*.

Note that the indicated motor mechanisms are intended to include "… all of the brain areas that contribute to overt behavioural response patterns."[109] This extension is truly extensive. Melzack and Wall note that the relevant motor mechanisms "extend throughout the whole of the central nervous system" and "must be at least as complex as … the input systems we have primarily dealt with."[110]

Melzack and Wall take their own gate-control mechanism to be well established, but in *The Challenge of Pain* they also claim that explaining the wide range of pain phenomena requires supplementing it with at least two other *independent* mechanisms:

> Each of the five stages that made up the gate-control mechanism has been supported by subsequent work. However, they are not sufficient to explain some of the basic facts about pain ... Two completely different mechanisms were later discovered which have no relation to the gate-control but which add to it and do not contradict it (176).

[109] Ibid., 193.
[110] Ibid., 193.

These two mechanisms are both mechanisms of prolonged pain: one impulse-triggered and one transport-controlled (or central).

The additional impulse-triggered mechanism was needed, most centrally, to explain inflammatory pain. Injury often results not only in immediate acute pain, but soreness, tenderness, and difficulty using the injured body parts. Aδ-fibres have slower conduction velocities than C-fibres, and as inflammatory pain develops later than acute pain, specificity theorists had posited that "slow pain" resulted just from Aδ-fibre input. Inflammatory pain, however, was found to crucially include transmissions from C-fibres. Further complexities arise because different types of tissue (joints, periarticular tissue, viscera muscle, deep fascia, and skin) show decreased ability to produce this long-term effect[111] Note the oddity in saying that any of your *organs* are sore. Melzack and Wall originally thought it could somehow be accounted for by gate-control, but it became clear that the explanatory mechanism would need to be centrally controlled. It turns out that C fibres hold a key to central changes because of their peptide contents, which change as a result of the kind of tissue they innervate. Much is still sketchy, but "[a] picture is beginning to emerge of a second system of transmission control in the dorsal horn which is dependent on peptides."[112] This mechanism, note again, is independent of the gate-control mechanism, and is needed to explain inflammatory pain (at least)—which the gate-control mechanism cannot.

The additional transport-controlled mechanism was initially posited to explain (other) kinds of centrally maintained pain, especially neuropathic pain, so far unexplained. The posited mechanism involves the transportation of chemicals other than the peptides released during C-

[111] Ibid., 178.
[112] Ibid., 180.

fibre stimulation discussed above. Melzack and Wall sum up this third mechanism and the rationale for accepting it in five main points:

1) When a peripheral nerve is cut, a cascade of changes sweeps centrally during the ensuing days and weeks, which alters the chemistry and physiology of the dorsal root ganglion cells, the motor neurons and the central terminals of the sensory fibres.
2) Changes in the afferents induce changes in the spinal cord, which include a reduction of inhibitions, a spread of receptive fields and an increase in excitability.
3) The delayed peripheral and central changes are not produced by nerve impulses since they are not altered by prolonged impulse blockade or by lesions, such as crush injury, which imitate the changes of afferent barrage produced by nerve section.
4) The time course and other properties of the central changes support the proposal that they are produced by changes in chemicals transported within the axons of sensory fibres. Since most of the changes can be produced by the specific C fibre poison, capsaicin, it is believed that the unmyelinated afferents play a particularly important role.
5) The chemicals and their messages remain unknown although nerve growth factor is a prominent candidate (*The Challenge of Pain,* 181).

In sum, Melzack and Wall recognize at least three independent pain mechanisms. The additional transport-controlled mechanism operates in parallel with the gate-control mechanism. The additional impulse-controlled mechanism is kicked off by the same impulses as the gate-control mechanism, but behaves independently.

Despite any claims made by their disciples, Melzack and Wall are clear that though they think that they have identified *a* mechanism implicated in many pain experiences, and not *the* pain mechanism. This is not surprising. The claim that there *is* a single pain mechanism is antithetical to the spirit of their quarrel with specificity theory; they not only reject that there is a single mechanism specific for pain, but that any of the mechanisms implicated in pain have components specific for pain processing. Throughout their work they deny that there are any receptors, cells, pathways or cortical centers that are properly understood as *pain* receptors, *pain*

cells, *pain* pathways, or *pain* centres. For one example of a passage characteristically antagonistic to such ideas:

> …central cells that respond exclusively or maximally to noxious stimuli are not 'pain cells.' There is no evidence to suggest that they are more important for pain perception and response than all the remaining somaesthetic cells that signal characteristic firing patterns about multiple properties of the stimulus, including noxious intensity. The view that only the cells that respond exclusively to noxious stimuli subserve pain and that the outputs of all other cells are no more than background noise is purely a psychological assumption and has no physiological basis (*The Challenge of Pain,* 157).

Melzack and Wall take themselves to have identified *a* mechanism for pain. That is certain. Their denial that it is *the* pain mechanism, or that it has components specific to pain processing, is equally certain.

Correspondingly, Melzack and Wall deny that they, or anyone else, have provided an adequate *definition* of pain. They repeatedly and explicitly deny that pain is any single type of experience; e.g., "… the term 'pain' is a broad label for countless different experiences."[113] They argue that the diversity defies definition:

> The diversity of pain experiences explains why it has been impossible, so far, to achieve a satisfactory definition of pain. The word 'pain' represents a *category* of experiences, signifying a multitude of different, unique experiences having different causes, and characterized by different qualities varying along a number of sensory, affective and evaluative dimensions.
>
> At present, we must be content with guidelines *toward* a definition rather than a definition itself. Too much remains to be learned about pain mechanisms before we can define pain with precision (*The Challenge of Pain,* 46).

Though they imply that they think we are headed towards a definition, Melzack and Wall's gate-control theory should nonetheless not be understood as providing one.

---

[113] Ibid., 156.

Despite not providing a definition or a single, underlying mechanism that explains pain, the gate-control theory revolutionized pain science. Prior to the theory, pain was taken to be exclusively determined by particular receptors, spinal cells, pathways, and brain centers. Pain science and the diagnostics it spawned were vested in a straitjacket that left numerous sufferers without relief: those who did not respond to standard treatments targeted at "the pain mechanism," were generally simply considered crazy. In the wake of gate-control theory, things have changed. The extensive influence of the theory scientific is less a result of its specification of what pain *is*, than of what it can *include*. Not so much a result of identifying *the* pain mechanism, but of identifying *multiple*, implicated mechanisms. Including *numerous* mechanisms as implicated in pain processing has advanced understanding and improved treatments:

> The gate-control theory of pain has provided, in large part, the conceptual background—the foundation—for new approaches to pain. The theory argues that pain does not have a single cause and is not even a single entity. There are multiple, interacting physiological and psychological mechanisms, and a rational approach to pain control requires multiple approaches that converge to produce a reduction in pain. (*The Challenge of Pain,* 273.)

Setting aside the *accuracy* of the model, the relevant moral of the story is that gate-control theory does not identify a single, underlying mechanism that explains pain. Indeed, the theory's authors repeatedly deny that there *is* a single mechanism or that pain is a single type of experience. The many new treatments are a testament to the better justice this complex account of pain does to the phenomena than any of its predecessors. Far from rendering pain a natural kind, however, this complex account straightforwardly denies that pain is explained by a single underlying mechanism.

3. Price's Somatic Perception Theory

Though Price's somatic perception theory of pain is heavily influenced by gate-control theory, it is worthy of a brief, separate discussion. As is standard, he follows Melzack and Wall in holding that pain has multiple dimensions that are genuine components and not mere reactions to pain: sensory, affective and cognitive. Thus his account is also complex. Departing from them, however, Price offers a definition of pain:

> The definition I propose [of pain] is a somatic perception containing (1) a bodily sensation with qualities like those reported during tissue-damaging stimulation, (2) an experienced threat associated with this sensation, and (3) a feeling of unpleasantness or other negative emotion based on this experienced threat. (*Psychological Mechanisms of Pain and Analgesia,* 2.)[114]

Price explicitly claims that these elements are individually necessary and jointly sufficient for pain.[115] The relevant sensation is held to arise via a distinct submodality of somatosensation and inherently disposes one to experience the feeling of threat and unpleasantness necessary for pain. These feelings are taken to be informed by the meaning each particular individual attaches to both the sensation and the threat itself. This meaning evaluation is cognitive and, on Price's view, a necessary, integral component of every token pain experience.[116]

Price often describes the three dimensions of pain he identifies as occurring in three stages: the initial sensation, followed by two distinct affective stages both causally affected by the particular individual's cognitive evaluation. Both affective stages are the "… end product of multiple contributing processes, including the pain sensation itself, arousal, autonomic and

---

[114] Price sometimes presents these elements in somewhat varied language. In some places, for instance, Price elaborates 'an experienced threat' as "…a meaning of threat to the body or self associated with this type of sensation." So too, the first condition, that the sensation is like those reported during 'tissue-damaging stimulation,' is usually weakened to sensations of variations of stimulations that *would* cause tissue damage *if* maintained over time.
[115] Price, *Psychological Mechanisms of Pain and Analgesia*, 8.
[116] Ibid, 7. See chapter 3 for a further elaboration.

somatomotor activation, and finally and most critically, cognitive appraisal."[117] The serial nature of these stages allows Price to conclude, as is intuitive, that pain really is unpleasant (stage 2 and 3) *because* it hurts (stage 1).[118]

Though Price makes much of the direction of causality from sensation to affect, it would be a mistake to construe him as presenting a purely linear model. As with gate-control theory, multiple descending pathways are embraced as necessary to explain pain's multiple dimensions. In addition, more than any of his predecessors, Price emphasizes the multiple *ascending* pathways implicated in pain processing. His theory is influential largely for emphasizing these ascending pathways and arguing that they are clues to understanding the psychology of pain—specifically, pain's cognitive-evaluative dimension as it is presented in his model. He offers these pathways in the following oft-presented schematic:

[117] Ibid., 43.
[118] Ibid., 51. Note that this intuitive conclusion only follows from Price's model if *hurt* is a sensation.

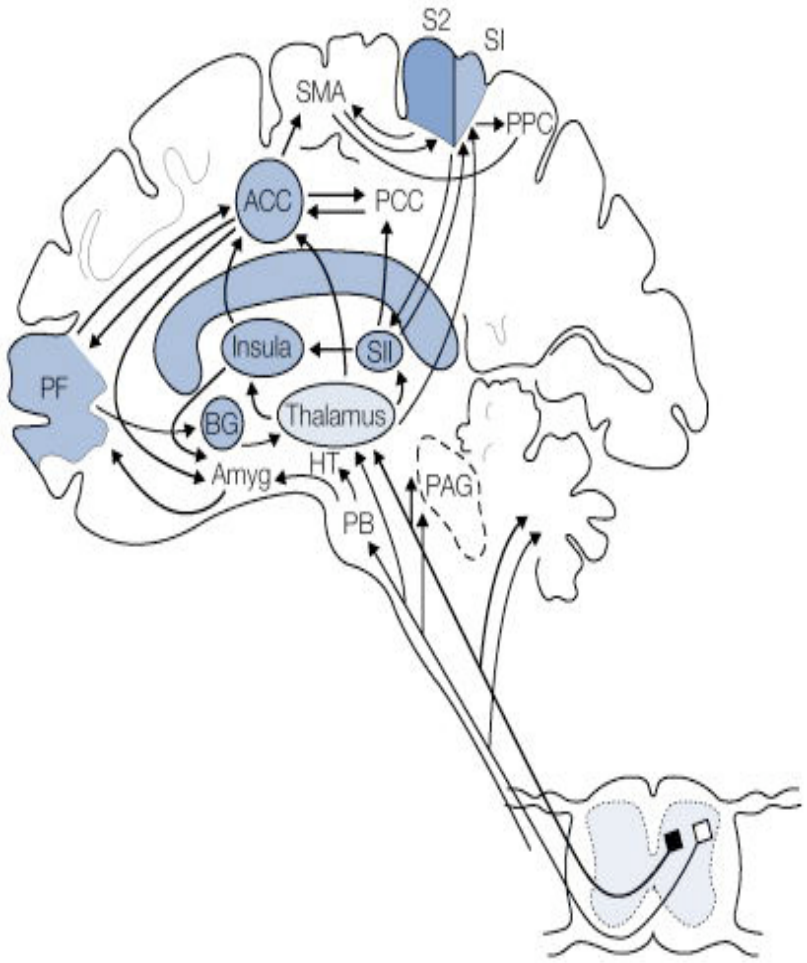

Figure 3.4:
Price's Pain Pathways
Taken from Price, 2000.

(ACC, Anterior cingulated cortex; Amyg, Amygdala; BG, basal ganglia; HT, hypothalamus; M1, primary motor cortex; PAG, periaqueductal grey, PB, parabrachial nucleus of the dorsolateral pons; PCC, posterior cingulate cortex; PF, prefrontal cortex; PPC, posterior parietal cortex; S1 and S2, first and second somatosensory cortical areas; SMA. supplementary motor area.)

Note that each of the above pathways transmit a variety of information and none are specific to pain processing.[119] The overall picture that emerges is one of multiple descending pathways (as emphasized by gate-control) and ascending pathways (depicted above)—each included to explain different sensory, affective, and cognitive dimensions of pain, but none specific for pain.

[119] Ibid., 99.

Price claims that each of these dimensions can be selectively influenced and, therefore, that assessing pain requires assessing each dimension.[120]

Moreover, Price believes that pain is necessarily a conscious experience that can only be known directly from the inside.[121] Because of this, Price holds that we must rely heavily on first-person reports to arrive at an accurate characterization of pain. Price thinks that this blocks a *reduction* of pain to any neurobiological or molecular mechanisms,[122] but that discovering distinct mechanisms that explain the distinct *dimensions* of pain experiences is an "ultimate goal" of pain science[123] To achieve this goal, pain scientists are exhorted to combine phenomenological reflection (trained first-person reporting) with more traditional third-personal methods. These methods, notice, are to be applied to *each* dimension of pain separately.[124] Like Melzack and Wall, Price explicitly and repeatedly claims that explanation of the *multiple* dimensions of pain requires identification of *multiple* mechanisms[125]

Setting aside the *accuracy* of Price's model or the utility of his methodology, the relevant moral of the story is that Price's somatic perception theory does not identify a single, underlying mechanism that explains pain. Though Price offers a definition of pain, it includes multiple components. Price's account emphasizes the multiplicity of pain's dimensions and takes a main goal of pain science to be an explanation of these various dimensions via postulations *multiple* mechanisms. Pain's multiple mechanisms are posited as operating both in serial and in parallel

---

[120] Ibid., 69.
[121] I have been intentionally putting off the question of whether pain is always conscious. While I seriously doubt it, whether pain is necessarily conscious is irrelevant for determining whether Price's account, even if accepted, renders pain a natural kind.
[122] Ibid., X.
[123] Ibid., 211.
[124] This methodology is further developed throughout Price's own work, and in Price and Aydede (2005). This work is important in its own right, but the details are irrelevant here.
[125] Ibid., e.g., ix, 99.

and none are specific to pain processing. Like gate-control theory then, this complex account also renders pain a non-natural kind: it straightforwardly denies that pain is explained by a single underlying mechanism.

4. Craig's Homeostatic Emotion Theory

Though not advocating a full return to specificity theory, Craig's theory of pain as a homeostatic emotion attempts to temper convergent views of pain—like the somatic perception and gate-control theories. Accordingly, he emphasizes the *specialization* of a specific pathway implicated in pain experiences: the lamina I spino-thalamo-cortical pathway. Note from the outset, however, that this pathway not only subserves pain, but homeostasis more generally—i.e., the body's internal regulation to maintain equilibrium. Other homeostatic emotions subserved by this pathway that Craig identifies include hunger, thirst, and itch.

A schematic of this system and its pathway is presented here:

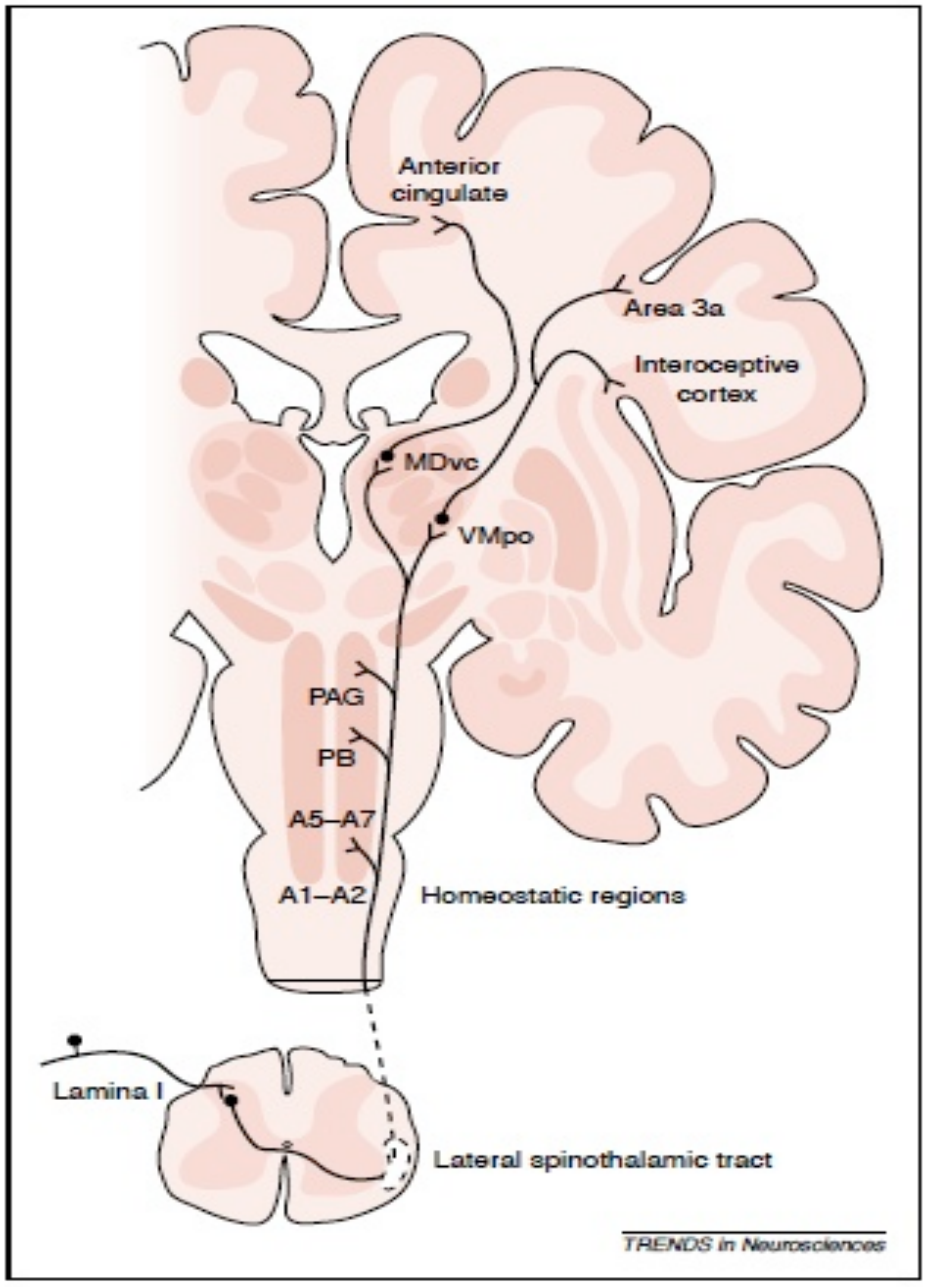

Figure 3.5: Craig's Homeostatic Model. Taken from Craig, 2003a.

Input from small-diameter fibres terminates in lamina I before being projected to the brain through the lateral spinothalamic tract. The identified projections are posited as explaining the sensory and motivational dimensions of pain—both of which are accepted by Craig as genuine

components of pain. In particular, the interoceptive cortex is taken to subserve pain's sensory dimension, and the motivational dimension is taken to be represented in the anterior cingulate cortex (ACC).[126] On Craig's view, every homeostatic emotion (e.g., itch and hunger) is comprised of a sensory and motivational component.[127] These components *together* constitute the homeostatic emotions, and it is for this reason that Craig's account is a complex account of pain.

Explaining pain, however, is not quite as simple as the schematic may be thought to imply. Craig holds that the experience of any homeostatic emotion, including pain, requires that information from the privileged pathway be integrated with other information in the forebrain.[128] This integration is not depicted above. The necessity of information integration makes sense if pain is a homeostatic emotion; information about the state of one's tissues must be coordinated with information about, for example, the position of one's limbs in order for one to perform an action that will achieve homeostasis using those limbs. The necessity of this non-modality-specific information shows that specificity theory is still rejected. Further complexities also arise for the subjective *awareness* of pain. For this, Craig holds that representations in the insular cortex must receive meta-representation in the anterior insula[129]—also not depicted. As Craig suggests that an experience is not *pain* unless we are aware of it, this additional step is essential to Craig's model.[130]

Craig's model then follows specificity theory by emphasizing specialized pathways, and follows convergent theories in emphasizing the integration of information from the activity in

---

[126] Craig, "Pain mechanisms: labeled lines versus convergence in central processing," 14.
[127] As discussed in Chapter 2, notice the similarity so far between Craig's account of pain as a homeostatic emotion and Clark's (2005) account of pain as a primitive feeling state.
[128] Ibid., e.g., 23.
[129] Ibid., e.g., 24.
[130] Ibid., 24.

multiple mechanisms. Craig therefore rightly takes himself to be offering a middle way in the ancient debate between specificity and convergence:

> The neural substrates underlying pain sensation in this model include both specific sensory channels (virtual 'labeled lines') [those depicted above] from lamina I and the convergent intensity-related pathways from lamina V [not depicted above]. Integration of these multiple ascending pathways in the brainstem and the forebrain is necessary for homeostatic control, for integrated perception, and for behavioral arousal. The essential conceptual difference between the prior convergent model and the more specific view described … stems from the fundamental recognition that the various feelings from the body represented in the STT [spinothalamic tract] are all aspects of the sense of the body's physiological condition, referred to as interoception, whereas the lemniscal fine touch and positional sense provide information about the relationship between the body and the external environment, that is, exteroception and proprioception. ("Pain mechanisms: labeled lines versus convergence in central processing," 23.)

Instead of contrasting specificity and convergence, Craig advocates emphasizing the distinction between interoception and exteroception. Craig maintains that the neurons in lamina I are specialized "physiologically, morphologically, and biochemically…",[131] and therefore laments the emphasis other convergent theorists placee on WDR cells for processing pain *or any other* homeostatic emotions. Homeostatic emotions do not result from exteroception, which is what Craig takes the WDRs subserve. Instead, they result from interoception, which is subserved by NS cells *and* a distinct class of cells posited by Craig called HPC cells (for heat, pinch, and cold respectively).[132] Note that HPC cells are held to be responsive to more than just noxious

---

[131] Ibid., 7.

[132] Craig addresses evidence cited by convergent theorists for the central role of WDRs in pain processing and the supposed insufficiency of the pathways he identifies to explain even pain's sensory dimension (Ibid., 14-22). He maintains that studies seeming to contradict his view result from mislabeling both the relevant cortical areas and cells in the dorsal column. I have neither the space nor the expertise to settle these debates here. For our purposes, it is enough to evaluate whether even *if* Craig's model is correct, it results in the naturalness of pain. See, however, Tommerdahl et al. (1996, 1998) for evidence that Craig's particular area of interest in the parietal cortex, area 3a, appears to respond preferentially to *all* thermal stimulation—not only noxious thermal stimulation. This research implies that the area is therefore also not best understood as

stimulation. Craig also laments that the crucial role of the insular cortex in homeostasis is under-emphasized in the dominant convergent models. It is, he argues, the specialized nature of both the lamina I cells and the insular cortex that allow the body to track the status of its internal state at all times. *Not*, notice, just when threatened. These specialized components are part of a hierarchical system in the body that subserves *all* homeostasis. Craig's central problem with other convergent theories is not that they fail to recognize a pain-specific system. It is, rather, that they inappropriately blur the lines between interoception and proprioception, thereby leaving us unable to appropriately identify the workings of the specialized *homeostatic* system. None of this is to deny his earlier claims that any homeostatic emotion, including pain, requires integration of information from multiple systems, or that awareness of the emotion requires a meta-representation of the integrated information.

In sum:

> … the neural representation of pain involves both specificity and integration. Pain as a homeostatic emotion is, in humans, both a specific interoceptive sensation and an integrated affective behavioral drive caused by a physiological imbalance that automatic (subconscious) homeostatic systems alone cannot rectify. ("Pain mechanisms: labeled lines versus convergence in central processing," 24.)

As with the other theories considered, let us set aside the *accuracy* of this model and see if it yields a natural kind for pain. It is more promising than the theories previously considered, because it focuses on a distinct system posited as explaining all homeostatic emotions—including pain. I do not deny that there is a homeostatic system or that reference to it in the generalizations of the relevant sciences may be useful. I lack the expertise to evaluate that claim. Even if there is a system for homeostasis, however, that does not show that there is a mechanism for pain.

---

nociceptive or pain specific—a conclusion that Craig himself needn't resist, given that his emphasis is on homeostasis and not pain alone.

Strikingly, there is nothing here about when the modeled homeostatic system results in pain as against any other homeostatic emotion. Craig is not even attempting to provide a mechanical model for pain in particular. A mechanical model for pain *as against a model of homeostatic control* cannot be garnered from focus on any subset of the components of the identified system: the receptors, the pathways, or the relevant cortical structures. This is to be expected since Craig explicitly follows convergent theorists in rejecting the idea that pain is explained by the workings of a specialized system. Craig, in fact, goes further than Price in denying that 'nociceptor' and 'nociception' are anything but heuristics. He emphasizes that small-diameter afferents and their receiving transmission cells not only process information about noxious stimuli, but argues that they are continually relating information about the body's status: the same mechanism is implicated when tissues are being damaged and when everything is just fine.

Presumably, Craig thinks pain is explained by *some* state or range of states of the system he identifies. We are simply not told which one(s). The following loose passage is the closest I can find Craig coming to distinguishing pain from any other homeostatic emotion:

> The basic homeostatic 'feelings' or modalities include temperature, itch, visceral distension, muscle ache, hunger thirst, 'air hunger' and sensual touch. All of these inherently generate an emotion that drives homeostatic behavior, and pain is no different. Pain normally originates from a physiological condition in the body that automatic (subconscious) homeostatic systems alone cannot rectify, and it comprises a sensation and a behavioral drive with reflexive autonomic adjustments. Pain can be either unpleasant (as usual) or pleasant (as when it relieves an intense itch). *The behavioral drive that we call pain* usually matches the intensity of the sensory input but it can vary under different conditions, and can become intolerable or, alternatively, disappear, just as hunger or thirst. ("A New View of Pain as a Homeostatic Emotion," 304, emphasis added.)

This is not a clear picture of when the homeostatic system is in a state, or range of states, appropriately deemed pain. Indeed, by speaking of "[t]he behavioral drive *that we call pain*" it

might be thought that even as Craig thinks 'nociception' and 'nociceptors' are simply heuristics, so is 'pain.' Craig even allows that pain may be pleasant depending on the other information being processed by the homeostatic system at the time.

In sum, while Craig's work focuses on pain as a homeostatic emotion, pain itself is not explicitly modeled and does not receive a mechanistic explanation. Moreover, no such explanation of pain can be extracted from Craig's impressive work. In line with his rejection of specificity theory, Craig's model seems, if anything, to support rejecting the very search for a pain-specific mechanism. The picture that emerges is a mechanistic explanation for *homeostasis*. Though homeostasis may be, pain is still not explained by the workings of an underlying mechanism. Even if Craig's model is accurate, it does not render pain a natural kind.

5. The Neuromatrix Theory of Pain

The neuromatrix theory of pain is the latest rage in pain science, but it is so new that the following brief treatment will no doubt be inadequate to conclusively dispel its future deliverance of pain as a natural kind. Despite its infancy, however, it is already thought by some to provide a pain model that can explain all types. My aim in this section is to show that it is *currently* unsuccessful in rendering pain a natural kind and to cast doubt on its prospects for doing so in the future. General considerations against *any* as yet unidentified mechanism explaining pain will be offered in the next chapter and are more decisive.

The initial presentation of the neuromatrix theory is traceable to Melzack (1989) in a discussion of phantom limb pain. Melzack explicitly posited the neuromatrix to explain pain phenomena that his earlier gate-control theory could not explain (e.g., phantom limb and chronic pain in particular). Moreover, the neuromatrix was not presented as a *pain* matrix, but as a

"body-self" neuromatrix: it is a model of the unified perception of the state of one's entire body—pain being one possible occurrent state among at any given time.

On Melzack's version, the neuromatrix is a "…large, widespread network of neurons that consists of loops between the thalamus and cortex as well as between the cortex and the limbic system."[133] Though these loops carry separate streams of neural information, they also regularly converge and synthesize. The synthesis of repeated patterns of processing is called a neurosignature. This signature "… is imparted on all nerve impulse patterns that flow through [the neuromatrix]."[134] Different neurosignatures are supposed to represent different overall, unified bodily states. Melzack posits the different patterns activated across the body-self neuromatrix encoded in the signature as explaining the different qualities experienced as happening to the body at different times. The *unified* experience of an embodied self is then explained by a projection of neural information stamped with the neurosignature to other areas of the brain. In the model, the areas receiving this projection are labeled as the "sentient neural hub." Here the stamped information is, somewhat mysteriously, "converted" into awareness. Projection of the neurosignature from the matrix to yet other neural networks is posited as explaining both movement (e.g., spinal reflexes) and the experience of movement.[135] Furthermore, though the spatial distribution of the matrix's loops is genetically determined, both the distribution and the patterns of information processing carried out across the loops become further individualized in response to sensory inputs throughout each individual's life. Each neurosignature is therefore highly idiosyncratic, even in the face of identical stimulation.[136] Accordingly, the theory's central thesis is that projected neural information "stamped" with my neurosignature will be

---

[133] Melzack, "Pain and the Neuromatrix in the Brain," 1379.
[134] Ibid., 1379.
[135] Ibid., 1379.
[136] Ibid., 1379-1381.

experienced as *mine*, constituting my subjective experience of a unified, embodied self. This is certainly intended to include my pain experiences, but is by no means limited to them.

Applied to the experience of pain in particular, Melzack (2001) thus sums up the theory:

> The neuromatrix theory of pain proposes that the neurosignature for pain experience is determined by the synaptic architecture of the neuromatrix, which is produced by genetic and sensory influences. The neurosignature pattern is also modulated by sensory inputs and cognitive events, such as psychological stress. It may also occur because stressors … act on stress-regulation systems … thereby contributing to the neurosignature patterns … The neuromatrix theory guides us … toward the concept of pain as a multidimensional experience produced by multiple influences. These influences range from the existing synaptic architecture of the neuromatrix—which is determined by the genetic and sensory factors—to influences from within the body and from other areas in the brain (1381).

Note the wide range of input. Melzack sometimes writes as if not only all sensory information, but all cognitive processing is causally relevant to the resultant neurosignature at any given moment of processing.

The representation of pain as a particular neurosignature (or, more likely given the sensitivity of the matrix, set of slightly variant signatures) is not intended to be at odds with the sensory, affective and motivational dimensions of pain Melzack earlier advocated. All of these components are taken to be captured by the activity of the neuromatrix and thereby encoded in the neurosignature(s) for pain. This multidimensionality is clear in Melzack's presented

schematic of the neuromatrix for pain experience:

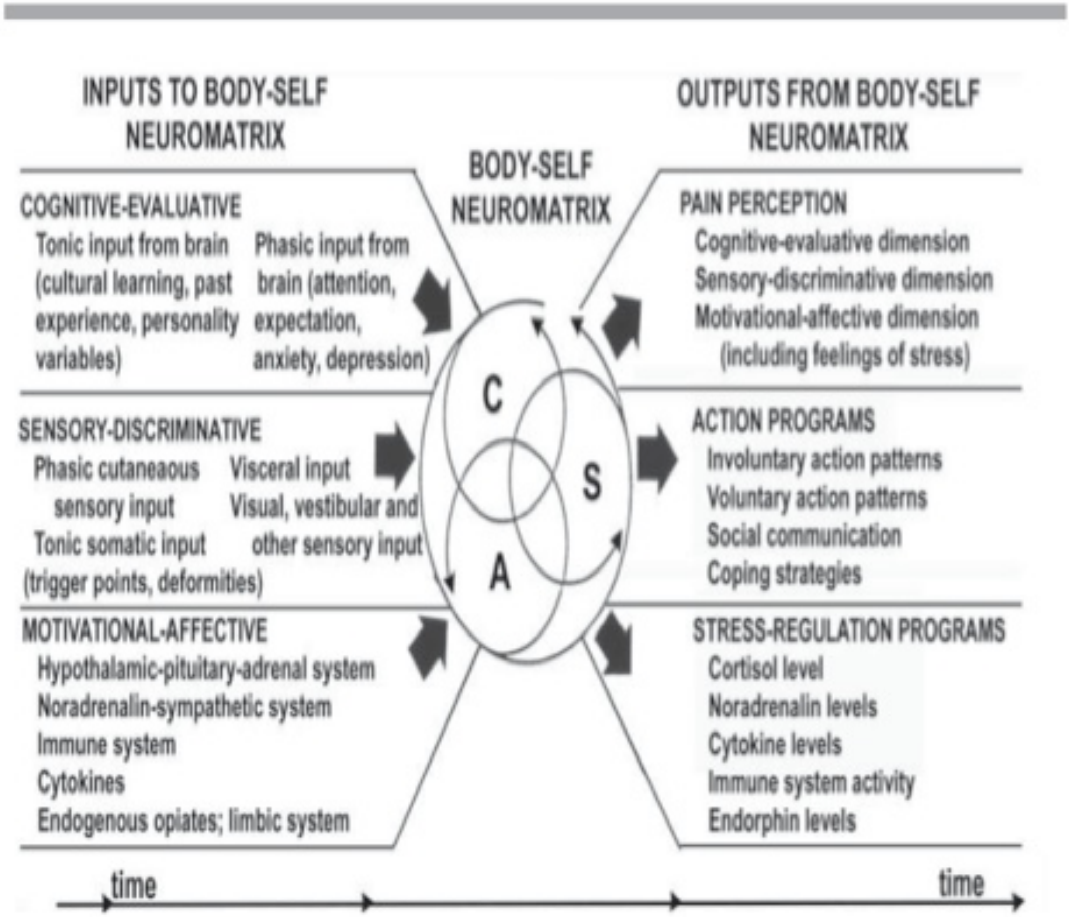

Figure 3.6: Melzack's Pain Neuromatrix. Taken from Melzack, 2001.

The sensory, affective and cognitive dimensions of pain ('A', 'C' and 'S' respectively) are now rendered "cognitive neuromodules," and are depicted as components in the body-self neuromatrix model of pain. Interactions between these components are then presumed crucial for determining the strucure of the experience. According to Melzack, every unified experience of the body-self will contain sensory, affective and cognitive components. Whether an experience is

a *pain* experience depends on the particular, individualized neurosignature encoded by the neuromatrix on the neural information flowing through it at any given time.

It is hard to evaluate whether, even granting the correctness of the theory, it yields a natural kind for pain; as with Craig's homeostatic system, there is no attempt made by Melzack to provide a model for pain in particular. Which neurosignature(s) explain(s) pain experiences? The slew of inputs, outputs and the standard three components presented in the schematic above are not specific enough to provide much explanatory power. The neuromatrix was explicitly posited to explain pain experiences that could not be explained by any mechanism whose activation relies on stimulation—like chronic pain. It is one thing, however, to say that these types of pains are not explained *merely* by modulation of nociceptive input. It is quite another to say what they *are* explained by. Without more detail, evaluating Melzack's proposal is difficult.

Though Melzack uses the expression 'the neuromatrix theory of pain' in the quoted passage, he nowhere refers to the neuromatrix as the *pain* neuromatrix. Such reference is a relatively recent addition to the development of the body-self neuromatrix. Pain researchers currently disagree about the relationship between the neuromatrix and pain experiences—Melzack's own view being one among many. On one end of the spectrum are respected pain researchers who consider the matrix to have specialized components for pain (e.g., Rainville 2002). Equally respected others on the opposite end hold that particular activations across the matrix merely *contribute* pain experiences (e.g., Tracey and Mantyh, 2007).

Though some researchers have taken to referring to the neuromatrix as *the pain* neuromatrix, recent evidence shows that this is likely inappropriate.[137] Mouraux et al. (2011) establish, conclusively to my mind, that when saliency is controlled, activity in the neuromatrix

[137] Many thanks to Rachel Corns for bringing this research to my attention.

is somatosensory specific but not nociceptive specific. They conclude that: "… the 'pain matrix', can also respond to any salient or behaviorally relevant stimulus, regardless of whether it is nociceptive in nature."[138] Similarly, Legrain et al. (2011) conclude that responses in the neuromatrix reflect the body's response to all salient sensory events—including those perceived as threatening *prior* to any experience the subject would report as pain (e.g., an approaching bee). They list and briefly describe a number of recent studies that support rejecting the claim that pain experience is represented by a specific pattern of activation within the matrix:

> Very recently, several studies have shown that this brain network cannot be reduced to a mere cortical "*representation"* of pain. Indeed, these studies have shown that the activity of the so-called 'pain matrix' (i) can clearly be dissociated from the perception of pain intensity … (ii) is strongly influenced by factors independent of the intensity of the nociceptive stimulus … and (iii) can be evoked by non-nociceptive and non-painful stimuli… Importantly, these experimental observations do not question the involvement of the cortical activity in the emergence of pain. Rather, they question the notion that the cortical activity involved in the generation of pain is necessarily and specifically reflected in the 'pain matrix' ("The pain matrix reloaded. A salience detection system for the body," 112).

These experimental observations and the conclusions of the these two studies may eventually be overturned. Again, the theory is quite new.

The theory is so new, in fact, that I submit that we are not yet presented with a mechanistic model for pain *at all*. Notice that even *if* pain experience was well correlated with a particular pattern of activity in the neuromatrix; such a correlation would provide only a description of the phenomenon. A description of the phenomenon is not alone explanatory. As emphasized throughout the mechanism literature and discussed in chapter 1: To explain, a mechanical model must explain how the identified activities and organization of the mechanism's components determine the phenomenon of interest. Even if Melzack's schematic is

[138] Mouraux et al., ""A multisensory investigation of the functional significance of the 'pain matrix',"2247.

accurate, we are a long way from such an explanatory model. The recent work by Mouraux et al. (2011) and Legrain et al. (2011) and those they canvas, suggests that even a phenomenal description of pain is not likely to be in the offing. If the activities of the neuromatrix do not even provide an accurate description of the phenomenon, that tells strongly against their eventually yielding an explanatory, mechanical model. In sum, even if Melzack's neuromatrix model is accurate as so far developed, it has not been thereby rendered pain a natural kind, and the recent research shows that prospects for its doing so are dim.

6. The International Association for the Study of Pain's Definition of Pain

The definition of pain offered by the International Association of Pain (IASP) was composed by prominent pain scientists and is often appealed to as if it were a scientific theory of pain. In this section I argue that that is not so. The IASP definition neither provides, nor intends to provide, an understanding of pain useful for the generalizations of any relevant science.

The IASP published its first definitions of pain terms in 1979. A committee was subsequently commissioned by the association to review the definition, leading to yet further reviews. The most recent approval of the definitions was in 2011. Numerous researchers, including Wall of gate-control fame, have contributed. The editor and final arbiter, however, has always been Harold Merskey. It is noteworthy that only minor revisions have ever been made to the original definition.

The oft-quoted definition currently reads: "An unpleasant sensory and emotional experience associated with actual or potential tissue damage or described in terms of such

damage."[139] Note that both sensation *and* emotion are explicitly included as components of pain and associated, but *only* associated, with tissue damage.

Getting at least as much attention in the literature as the definition, is its lengthy appended note:

> Note: The inability to communicate verbally does not negate the possibility that an individual is experiencing pain and is in need of appropriate pain-relieving treatment. Pain is always subjective. Each individual learns the application of the word through experiences related to injury in early life. Biologists recognize that those stimuli which cause pain are liable to damage tissue. Accordingly, pain is that experience we associate with actual or potential tissue damage. It is unquestionably a sensation in a part or parts of the body, but it is also always unpleasant and therefore also an emotional experience. Experiences which resemble pain but are not unpleasant, e.g., pricking, should not be called pain. Unpleasant abnormal experiences (dysesthesias) may also be pain but are not necessarily so because, subjectively, they may not have the usual sensory qualities of pain.
>
> Many people report pain in the absence of tissue damage or any likely pathophysiological cause; usually this happens for psychological reasons. There is usually no way to distinguish their experience from that due to tissue damage if we take the subjective report. If they regard their experience as pain and if they report it in the same ways as pain caused by tissue damage, it should be accepted as pain. This definition avoids tying pain to the stimulus. Activity induced in the nociceptor and nociceptive pathways by a noxious stimulus is not pain, which is always a psychological state, even though we may well appreciate that pain most often has a proximate physical cause.

Pain is defined subjectively here, and the definition indirectly builds in that a person is not only incorrigible, but infallible, about whether they are in pain. If pain is sincerely reported, then pain is occurring—whatever physical examination might reveal. Indeed, especially in its last sentence, the definition appears to contrast the physical and the psychological altogether. Having contrasted these, the IASP attempts to unify all pain experiences by focusing on the psychological alone instead of decreeing some pains physical and some psychological.

---

[139] http://www.iasppain.org/AM/Template.cfm?Section=Pain_Definitions&Template=/CM/HTMLDisplay.cfm&ContentID=1728#Pain

According to the IASP then, pain is never constituted by a physical process, is always psychological, and only sometimes has a physical cause.

While the definition has some intrinsic adequacy problems,[140] were it accurate, it would still not render pain a natural kind; it does not even attempt an explanation of pain, much less a mechanistic one. It explicitly distances itself from any theory that requires pain to be tied to, much less determined or explained by, any physical process. Despite how this definition is often employed in the philosophical and scientific literature, I argue that the definition was never intended fur use in scientific generalizations; it was never intended to be a scientific definition of the nature of pain.

In 1996, Anand and Craig attacked the definition on the grounds that it inappropriately makes all attributions of pain to creatures unable to report their states illegitimate. For present purposes, the charge itself is less interesting than Merskey's and Wall's defenses against it. These defenses help illuminate the intended relationship between the offered definition and scientific inquiry.

Merskey is clear that the definition is not intended to be "… a statement of absolute truth but a statement of the way in which we employ a term."[141] Note that the relevant "we" here is not the scientist, but the layman. He admits that the definition has contributed to inadequate care of infants and other pre or non-verbal persons.[142] He claims, however, that this is not his problem.

---

[140] Valerie Hardcastle, for instance, notes that if the IASP definition of pain is correct and pain is purely subjective then "… pains would be deeply peculiar. They would be the only perceptual experiences we have that are normal (that is, normal functioning creatures have them), natural (that is, not the product of tweaking something in a laboratory, but occurring in the wild, as it were), commonplace (we have them all the time), believed to be giving us information about the external environment (external to the brain, that is), but in fact do not" *The Myth of Pain*, (129-130).

[141] Merskey, "Response to Editorial: New Perspectives on the Definition of Pain," 209.

[142] Ibid., 209.

He maintains that his definition is correct if it accurately reflects common usage. He notes that we may, someday, be able to correlate pain (understood as the subjective experience described in the IASP definition) with "pain patterns"; i.e., pain behavior or neural activity. Nonetheless, he holds that we cannot yet make these correlations: he believes that we cannot correlate pain with *anything* observable. Merskey is explicit that we nonetheless ought to follow common usage in defining 'pain,':

> It is a misdirection of energy to attempt to change the meaning of a word that, so far as I can tell, has been used everywhere—or perhaps almost everywhere—in a particular way for thousands of years. Efforts should rather be directed to a steady improvement of the recognition of probable pain patterns and their predictive value in relation to treatment ("Response to Editorial: New Perspectives on the Definition of Pain, 209).

Merskey appears to be a bit muddled about definitions; he seems to take it for granted that defining 'pain', defining pain, making explicit common usage, and identifying associated descriptions of pain, are all one and the same task. This is surely not so. What I would like to emphasize, however, is that the central goal of the committee was to make common usage explicit.

Merskey explicitly distances the layman's understanding of pain, as he thinks he has captured it, from treatment. Treatment should focus on "pain patterns"—but he holds that none of these are correlated with anything 'pain' is commonly used to identify by the layman. We should attempt to treat, and research, those things we can observe—but Merskey thinks that none of these correspond to what the folk are talking about when they use the term 'pain.' Until we can correlate something observable with the thing that Merskey thinks common usage picks out, "… then pain remains a word which is both powerful, because of its restriction to the subjective

state, and limited for that very reason."[143] Putting aside some semantic muddling, Merskey's point here, again, is that though nothing observable corresponds to common usage, common usage is not thereby altered, nor should we attempt to alter it. Notice what a ludicrous suggestion this would be if Merskey was intending to give a *scientific* definition of pain or to identify the *nature* of pain.

Wall's own (1996) brief response to the challenge presented by Anand and Craig is similar and equally instructive. For treatment purposes, he decries the very question 'Is this organism in pain' in favor of the question 'Would you interfere with the integrity of an organism.' This response is meant to highlight that Wall thinks these are very different questions. Questions about treatment concern whether certain interventions ought be made—and Wall's answer turns on whether the intervention will benefit the organism. Whether an organism is in pain, however, is not a useful question for whether to take action or what kind of action should be taken. For an intervention, there are anatomical, biological and behavioral considerations. Wall believes that none of these correspond to the folk notion of pain. A sincere report of pain tells you if the person is in pain, and that alone, according to Wall, tells you nothing useful for treatment about the state of the organism. This is not to deny that *more specific descriptions* of the qualities, duration, and history of the pain are diagnostically useful. The point is that the binary question—pain or no—is diagnostically inert.

Indeed, the IASP's website emphasizes the limited use of its pain definition in their oft-ignored preamble. They note that the definition of pain offered is *not* intended as a "constraint on further developments," and that "[i]t is important to emphasize… that the terms have been developed for use in clinical practice rather than for experimental work, physiology or

---

[143] Ibid., 209.

anatomical purposes." This identified aim is in line with Merskey's defense of the definition as, and only as, an accurate reflection of common usage. Though the IASP's preamble is not clear about *how* the definition is to be employed in clinical practice, I submit that it is useful insofar as it enables doctors to better communicate with their patients. Doctors should be informed about common usage as most patients will, no doubt, employ the term that way. Whether or not this was the intended clinical use of the definition, that it is *not* intended for research purposes, even only to *constrain* them, makes clear that it does not present anything like a mechanistic explanation of pain or a scientific definition.

Anand and Craig's criticism of the definition appears to stem, in fact, from their dissatisfaction with its limited aim. In response to Merskey and Wall's defense, they write that:

> Whereas common usage of the word 'pain' has to be learnt from early experiences and has remained unchanged for thousands of years, the scientific definition of this phenomenon must go beyond the semantic use of a word to the core of the experience ("Reply to Letters to the Editor from Merskey & Wall, 210).

Setting aside concerns about the semantic methodology enshrined in this passage, insofar as Anand and Craig are claiming that our scientific understanding of pain must go *beyond* common usage, I think they are quite right. A "scientific definition" of pain should provide a referent useful for science's generalizations. Common usage, certainly, does not always provide this. They err, however, if they think that Merskey and Wall disagree with them. The IASP definition was just never intended as a "scientific definition"—it was explicitly offered under the proviso that it was unhelpful for research.[144] It was not intended to elucidate the nature of pain, correlate with any physical process, constrain future research, or be used in any context outside the

[144] Philosophers who regularly appeal to this definition would do well to note this!

clinical setting—where, I conjecture, it is useful *insofar* as it facilitates communication between doctor and patient.

Anand and Craig's cry for a new definition is understandable since the IASP definition has, in practice, been used and appealed to as if it *were* a scientific definition. Those who make these appeals, however, have not been attentive to the stated purpose of the definition offered by the IASP, or Merskey's direct defense of its limited use. The IASP definition does *not* render pain a natural kind—it does not identify a referent of the term for use in any scientific generalizations, much less identify a mechanism that explains the phenomenon.

7. Conclusion

In this chapter, I surveyed the 5 models of pain currently dominating pain research. None of these accounts offer a mechanistic explanation of pain; none identifies an underlying, homeostatic mechanism that explains the wide range of pain experiences. If these complex accounts of pain were all that we could ever hope to have, then we could confidently conclude that pain is not a natural kind. Of course this is not so. Despite the failures of these dominant accounts to render pain natural, it may still be hoped that some future model of pain will succeed.

In the next chapter, I turn to recent research on pain *types*. It is now well established that different mechanisms are implicated in different pain experiences. The claim that there is no underlying, homeostatic mechanism that will explain all pain is most directly supported by the fact that different pain experiences are explained by the activity of different, independent mechanisms. The research on pain types and the new mechanism-based approach to classifying pain, moreover, strongly supports the further conclusion that no type of pain is itself a natural

kind. At the end of chapter 4, I therefore conclude that neither pain, nor any type of pain, is a natural kind.

Chapter 4: Pain Types

1. Introduction

The claim that pain is not a natural kind is appropriately broken down into two: 1) Pain is not a natural kind; and, 2) No type of pain is itself a natural kind. In support of claim 1, the previous chapter argued that the dominant pain models currently on offer do not renders pain a natural kind. It may still be rightly objected, however, that those models' failures do not establish that no *future* model will succeed. I establish this and claim 2 in this chapter by examining what we already know about the mechanisms implicated in different types of pain experiences.

To this end, section 2 employs one standard categorization of pain into three main types: nociceptive, chronic, and inflammatory pain. The discussion focuses on the fact that these different types receive different mechanistic explanations: different mechanisms, with different components, are implicated in some types of pain and not others. This is most clearly evidenced by different *treatments* for different types of pain—a piece, even, of common knowledge. The blatantly different mechanical models and interventions for different types of pain strongly supports the conclusion that all pains are not determined by a single, underlying mechanism. Moreover, the known variation in determining mechanisms across pain types undermines any reason to think there will be a *future* explanation, of all pain by the workings of a single, underlying mechanism. At the end of section 2, I take claim 1 to be established; pain is not a natural kind.

Having thus established claim 1, in section 3 I argue for claim 2—no type of pain is itself a natural kind. Recent work in pain science attempts to provide mechanism-based classifications and explanations of pain types, and the admitted failures thus far are instructive. What they are finding, in short, is that each token pain is the result of the convergent activation of *multiple*

mechanisms—genetic, chemical, biological and social. The central reason no mechanistic explanation even for types of pain is forthcoming is that the convergent activity of these multiple mechanisms is idiosyncratic. the explanation of each token pain, moreover, is not only idiosyncratic for individuals, but to the same individual at different times. This idiosyncrasy is crucial for explaining pain tokens, but also blocks useful *generalizations* about pain "types."[145] I conclude that different types of pain do not submit to a mechanistic explanation, and that no discovery of yet additional mechanisms implicated in pain experiences will help. Claim 2 will thus be established; no type of pain is itself a natural kind.

An objection rightly prompted by the current and previous chapter is that 'pain' is being employed to do some useful scientific work. This utility appears to be at odds with the scientific utility criterion presented in chapter 1; if 'pain' is being usefully employed in scientific inquiry, then isn't pain a natural kind according to my own criterion? The objection is serious, and I deal with it in section 4. In brief, the refutation is that it is pain scientists' very inquiries into pain that has made clear that the term 'pain' and terms for pain "types" are no longer useful. This is analogous to the fruitful inquiry of physicists into luminiferous aether and of doctors and biologists into the humours—inquiries that illuminated that the terms 'aether' and 'humours' were not useful. Paradoxically put: Our pain science has surpassed 'pain.' The term is not usefully referred to in the generalizations of the relevant science and is therefore, consistent with the scientific utility criterion, not a natural kind.

[145] My use of scare quotes throughout marks only what I think is the artificiality of the constructed types. One can, of course, construct sets out of any tokens one wishes, but that does not thereby a useful type make. I use the scare quotes whenever I think the artificiality of these constructions bears emphasizing.

2. Different Types of Pain Do Not Constitute a Single, Natural Kind

Numerous distinctions are employed to distinguish among different types of pain for research and diagnostics. The simplest is the distinction between acute and chronic pain; i.e., pain of brief and long duration. It is now well known that acute pain and chronic pain implicate different mechanisms and, accordingly, are appropriately treated differently. A discussion limited to only these two types of pain would be enough to show that they do not constitute a natural kind. Rather than focus on this simplistic dichotomy, however, I'll focus on the most standard classification of pain types: the three broad types of nociceptive, inflammatory, and chronic pain. What is known about these broad types shows that pain is not a natural kind.

As a preliminary, note that many pain researchers and clinicians categorize pain into more types than the three I focus on. These additional classifications are usually more fine-grained, falling within one of the three identified categories. Obviously, if the three types considered here do not constitute a natural kind, the same will go for the more fine-grained types they include. Alternatively and less commonly, some types of pain are taken to implicate mechanisms such that they cross-cut the broad types I focus on. In "Bone Cancer Pain: From Model to Mechanism to Therapy," for instance, Luger et al. (2005) argue that bone cancer pain implicates the activity of unique mechanisms with components similar to both inflammatory and neuropathic pain, but different enough that bone-cancer pain is inappropriately classified as either. Additional pain types like these are of no concern for my purposes; if the pain types focused on here do not constitute a natural kind, adding yet *more* types will not help. Whether or not *any* type itself (e.g., bone-cancer pain) constitutes a natural kind is a more difficult question taken up in the next section.

Scholz and Woolf (2002) discuss nociceptive, inflammatory, and neuropathic pain in taking up the question "Can We Conquer Pain." The article explicitly emphasizes that different types of pain are explained by different mechanisms. That we *call* all of these types *pain,* does not entail that they should not be categorized distinctly:

> *[a]lthough* we use a single word to describe any feeling that is unpleasant and hurts, this does not mean that pain is a monolithic entity. … Pain the sensation can be further split into distinct categories" (1062).

The authors' stated reason for splitting pain into distinct categories is mechanistic:

> Nociceptive, inflammatory and neuropathic pains result from diverse mechanisms. Some of these mechanisms are unique to one pain condition; others are present in multiple clinical syndromes, or may be expressed at different times during the natural history of a syndrome. In some patients, a single mechanism may produce their pain; in others, multiple mechanisms may contribute (1062).

Standard mechanical models of each pain types is provided by the authors and re-presented below.

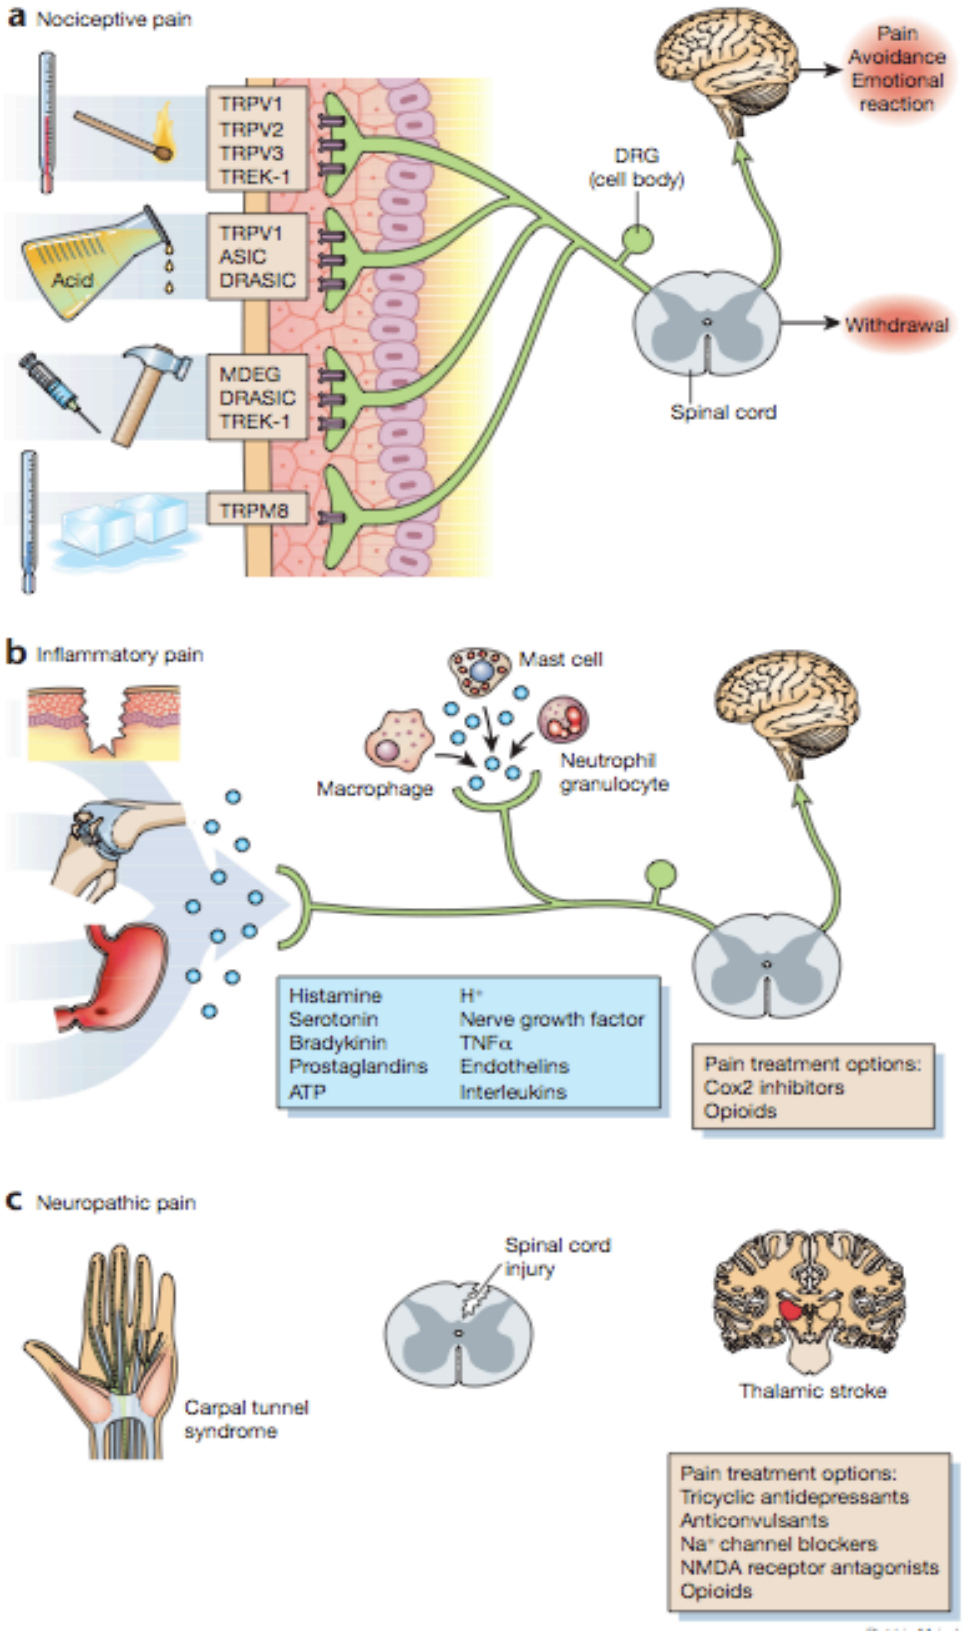

Figure 4.1: Three Broad Pain Types. Taken from Scholz and Woolf, 2002.

The details of these models are not relevant, but it is important to see that the different purported types of pain have different mechanical models. Though all schematically include nociceptors, the dorsal horn, and the brain, each also presents different activities in these regions and additional (different) components. If the models presented here are accurate, they are enough to show that different pain types implicate different mechanisms.

For doubters maintaining that there just *must* be some unified pain mechanism, and we should keep trying to find it, the most effective line of refutation may emphasizing treatment. Note the pharmacological treatment options offered on the right side of the models above. Notice that none are even provided for nociceptive pain—reflecting that pharmacological treatment of nociceptive pain is usually inappropriate. Consider the absurdity, for instance, of taking an aspirin for a paper cut. That different pain types require different treatment is common sense. With perhaps no, or a merely vague, understanding of the relevant mechanisms: you take Pepto-Bismol for stomach pain and Tylenol for a headache. *When* these different treatments are effective for different "types" of pain, it is because the activity of different mechanisms are involved.

These points about treatment are helpfully connected to the mechanistic literature canvassed in Chapter 1. In *Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience*, Craver repeatedly emphasizes that the *goal* of mechanistic explanation is manipulation and control. This claim is overblown: explanation is one thing and control is another. A scientist's goal may be to explain something that they know they lack the power to control, or that they think it immoral to even attempt to control. Consider, for instance, anthropological explanations, or the origins of the universe. These sorts of scientific inquiry seek mechanistic explanations, but not control or manipulation. In pain science, however, Craver's

point is well taken. The goal of explanation in pain science is manipulation and control; i.e., the goal is treatment. A good mechanical model of pain would facilitate the explanation and predictions necessary to effectively manipulate and control it. Scholz and Woolf (2002) are clear that the reason for identifying the involvement of different mechanisms in different types of pain is so that patients can receive adequate care.

Nociceptive pain, inflammatory pain, and neuropathic pain are not explained by a single underlying mechanism and, therefore, do not identify a unique target for treatment. As such, within pain *science*, pain should be understood as a heterogeneous phenomenon, because it is only so understood that adequate treatment is possible:

> The notion that there is a class of drug, a universal analgesic, that can intrinsically reduce all pain, is obsolete and has to be abandoned. Pain is heterogeneous in terms of etiological factors, mechanisms and temporal characteristics. Consequently, treatment must be targeted not at the general symptom, the pain, or its temporal properties, acute or chronic, but rather at the underlying neurobiological mechanisms responsible. ("Can We Conquer Pain, 1065.)

Recognition of different pain types and multiple pain mechanisms is standard among clinicians and pain researchers, but one may still demure that Scholz and Woolf have gone too far. It may be accepted that different treatments are appropriate for different kinds of pain, while still hoping for a universal analgesic. A universal analgesic, if discovered, would presumably work by acting upon some universal, underlying pain mechanism—showing that pain is a natural kind after all. A treatment effective on all pains would implicate an as yet unknown mechanism somehow underlying those currently known.

An oft-quoted review by Edwards et al (1999) is sometimes offered in support of the contention that there *is* a universal analgesic: aspirin. The study concluded that aspirin had significant analgesic effects on all postoperative pain—independently of pain model, duration, location or even the quality of the studies included in the review. Some of the authors' own

comments may suggest that aspirin is the long-sought panacea. Consider the following passage from the discussion:

> This review provides empirical evidence from aspirin that results will be the same for different acute pain states, different pain measurements, different group sizes, different quality of study design, and different study durations. This should help future reviewers who wish to study interventions in postoperative pain. It is important empirical support for the 'lumpers', who wish pragmatically to pool all admissible data, against a 'splitter' contention that pain is different at different sites and responds differently to analgesics. ("Oral aspirin in postoperative pain: a quantitative systematic review," 295.)

Concluding that aspirin is a universal analgesic on the basis of this review, however, is over-hasty. Notice first that the authors speak of "acute pain" presumably in contrast with chronic pain—which the review does not include. More important, notice that the studies included in the review were all explicitly on *postoperative pain*. Some distinct mechanisms are, indeed, active upon trauma to the body. Aspirin, as is now well known, works by blocking the inflammation that results from tissue damage:

> The site of action of these drugs [aspirin and other non-steroidal anti-inflammatory drugs or NSAIDS] appears to be entirely on the injured tissue itself and there is no convincing evidence that therapeutic doses have any effect directly on peripheral nerves or in the central nervous system. … When tissue is damaged, prostaglandins are synthesized and released into tissues where they play a part in triggering the three classic signs of inflammation… Aspirin blocks the synthesis of prostaglandins and therefore prevents the appearance of this crucial substance which announces that tissue is damaged. That is, it acts within the damaged tissue itself and not on the nervous system. (*The Challenge of Pain,* 200.)

This explains the unity found in aspirin treatment for post-operative pain, without jumping to the hasty conclusion that all pain, even all acute pain, is effectively treated with aspirin.[146] For all that is determined by this infamous review, aspirin may be appropriately employed for all pain involving tissue damage which may be granted without admitting that aspirin treats all *pain*. This

[146] For a more extended, and technical, discussion of why Edwards et al. (1999) do not establish that all acute pain is explained by the same mechanisms or appropriately treated identically, see Chapter 42 of *The Pain Textbook*.

is so because not all pain involves tissue damage—as illustrated in the mechanistic models above and reiterated throughout the previous two chapters. Even *if* the synthesis of prostaglandins is blocked even "acute pain" may be experienced. So too, notice that the drugs' efficacy is an average across what were admitted to be varied individual responses. As a result of the variation between pain and tissue damage, and the variation across individual responses, there are some pains that aspirin leaves completely untouched—as we already common-sensically know that it does.

The NSAIDs (aspirin being one) may be the best candidate for a universal analgesic, but they fail to deliver. Those who maintain that there just *must* be a universal analgesic and an underlying pain mechanism should, at this point, provide some reason for thinking there is any such thing.

A different objection to the claim that the different types of pain do not constitute a natural kind is the crucial role of pain reports in diagnostics. That role, however, further supports pains' heterogeneity. An illustrative example is the clinical use of the McGill Pain Questionnaire (MPQ). The questionnaire, introduced in chapter 2, is the most internationally respected tool for soliciting pain reports for diagnostic purposes. The long form of the MPQ is:

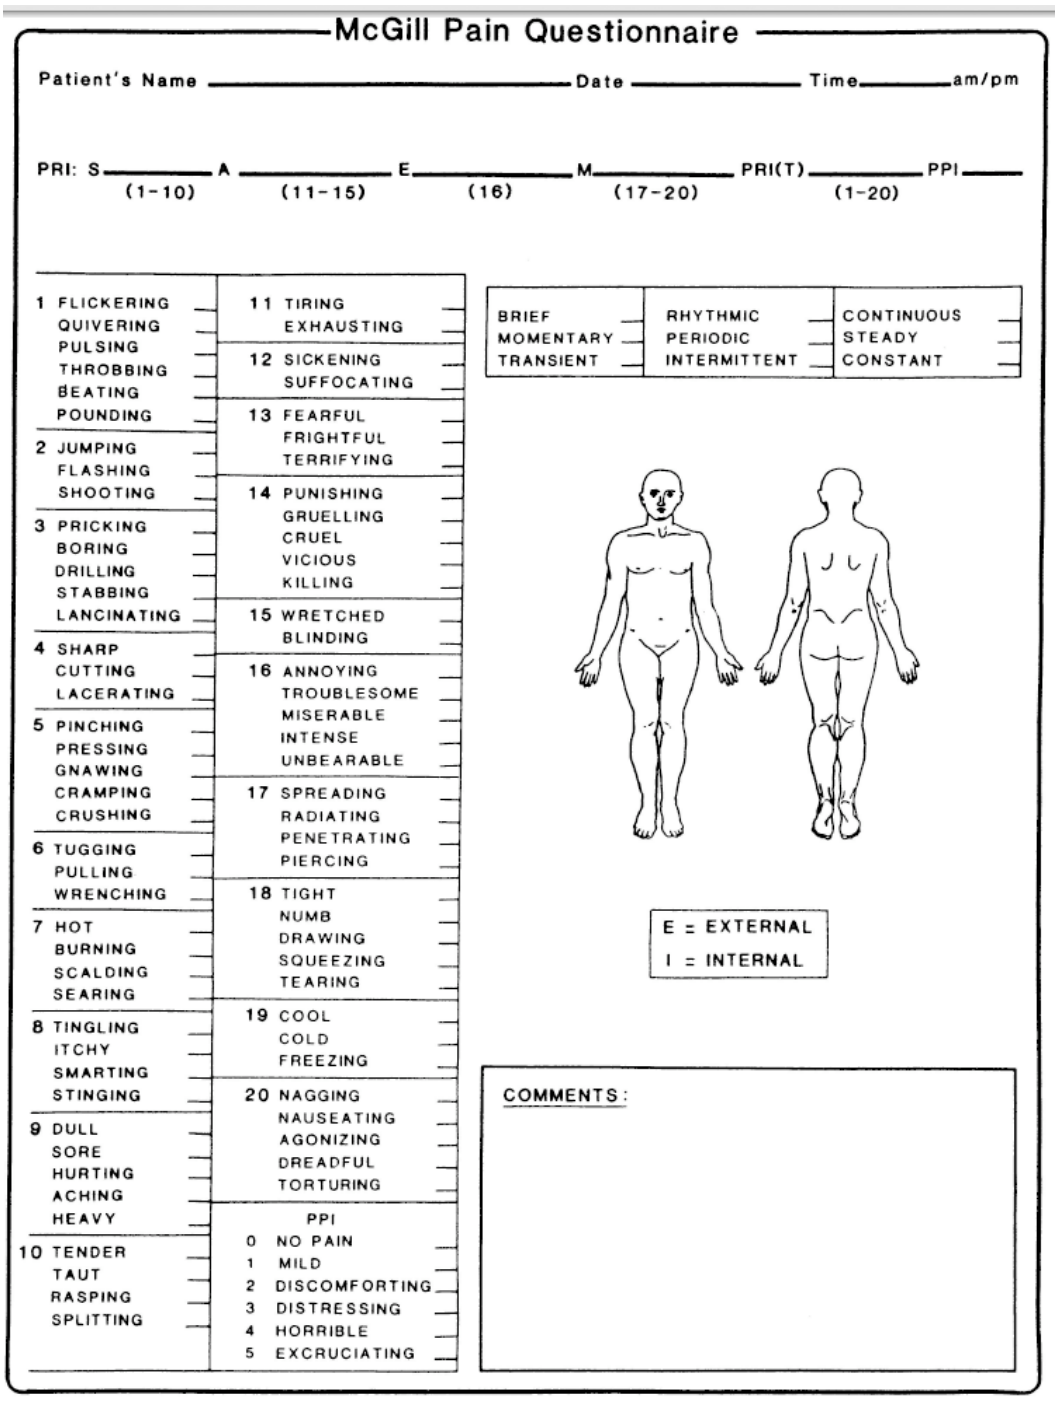

Figure 4.2: The McGill Pain Questionnaire.

The questionnaire's impact stems directly from its diagnostic utility. Its accuracy was tested in a study reported by Melzack and Wall (1988). Patients antecedently diagnosed with one of eight different pain syndromes described their pain using the questionnaire. The descriptions were fed into a computer with a diagnostic program and the results were impressive:

> The computer made a correct classification in seventy-seven per cent of the cases. When the sex of the patient and the location of the pain were also included, the classification was correct in one hundred percent of the cases. It is evident, then, that there are appreciable and quantifiable differences in the ways various types of pain are described, and that patients with the same disease or pain syndrome tend to use remarkably similar words to communicate what they feel (*The Challenge of Pain*, 43).

This success, however, should not be overestimated. Complications that lead into the next section arise by noting that diagnosis here is of a disease, pathology or syndrome—not an identification of a mechanism. While supporting the consistency in MPQ descriptions chosen by patients *diagnosed with the same syndrome*, the study is useless for establishing the legitimacy of the diagnostic criteria themselves. The computer gets it right because it has been programmed to associate particular constellations of descriptors with particular diagnoses. The *same* constellations of descriptors that the doctors employed to make the original diagnoses. Showing that the diagnoses based on these descriptors are useful for pain treatment requires evidence that the diagnoses provide appropriate *targets* for treatment. The recent work on mechanism-based classifications of pain discussed in the following section reveals that it is because diseases and pathologies are *not* well-correlated with mechanisms that these diagnoses do *not* provide effective pain-treatment targets. This is so even if the diagnoses are, trivially, correlated with reports.

Equally relevant for present purposes, notice that the general report 'I am in pain,' is diagnostically inert. No diagnosis can be provided on this basis. The utility of pain reports

depends entirely on the provision of further information. Consider your bizarreness reaction if a medical professional asked you if *it* hurt, and refused to provide any further specifications. Reporting pain, without further details about the token pain—its location, intensity, etiology, and other qualities—is useless. As I've been arguing, this is what we should expect since pain is mechanistically heterogeneous: a binary pain report does not identify an underlying mechanism that can be targeted for treatment.[147]

The different types of pain do not constitute a natural kind. As argued in chapter 1, if pain is a natural kind, then it should be possible, in principle, to arrive at a mechanistic explanation of it. As we saw in the last section, the dominating, complex models of pain fail to provide such an explanation. The more we investigate, the more we find that *different* pains implicate *different* mechanisms; even the three broad pain types discussed here implicate different mechanisms. Mechanistic differences are evidenced most clearly by focusing on treatment—the goal of pain science and any hoped-for mechanical model of pain. The general type 'pain' is diagnostically inert—even an inadequate diagnosis by pathology or disease requires distinguishing among types of pain experiences. There is no universal analgesic precisely because there is no pain mechanism. There is no single, underlying mechanism that explains all pains; pain is not a natural kind. We now turn to the question of whether different types of pain are themselves natural kinds.

3. No Type of Pain is itself a Natural Kind

As briefly discussed in chapter 1, until recently, pain research and diagnostics employed what is called the disease, or medical, model of pain. This model emphasizes an explanatory role for

[147] Recall Wall's similar point discussed in the previous chapter that the binary question "pain or no" is unhelpful for treatment purposes.

pathologies. A patient presenting with pain is to be diagnosed such that the presumed pathological origin of their pain is identified and, when possible, eliminated. The core assumption is that different types of diseases and pathologies explain different types of pain; to treat the pain one need only treat the pathology.

The model is intuitive and may even seem inevitable, but it has recently come under attack by many pain scientists. Practically speaking, the medical model has failed to facilitate adequate pain treatment. Recall the useful passage from Smart, et al (2008) that helpfully elaborates the traditional medical/disease model approach and problems with its assumptions:

> The medical/disease model refers to a tissue and pathology oriented approach towards the explanation of pain and dysfunction and functions on the premise that all pain has a dominant tissue or structural source. … The medical model is based on three broad assumptions:
> - There is a direct and proportional relationship between the amount of tissue damage and the level of pain experienced.
> - There is a linear relationship between the amount of tissue injury and the degree of any ensuing physical impairment or disability.
> - Treating the pathology resolves the pain and therefore the disability.
>
> … [I]nterpretations and classifications of pain based on the medical/disease model and its structure-oriented approach have been criticized since such approaches do not explain the complexity and variability of pain presentations observed clinically; where pain is observed to persist after healing, where patients' reports of pain in response to similar injuries differ greatly or where the magnitude of a pain report does not match the injury or pathology. Nor do they account for the variations in patients' outcomes in response to interventions directed towards its treatment since by the tenants of the traditional medical model treatment of the pathological process should be accompanied by relief of pain (2).

The moral of the story is that there are patients with the same pathologies and varying (including no) pain, similarly-reported pains and varying pathologies, and pains with seemingly *no* underlying pathology. Diseases and pathologies are structural or functional problems with an organism that need to be eliminated. Classification of *pain* by disease or pathology, however, does not adequately classify pain experiences; the two independently vary. Some pains are explained by an underlying disease or pathology, such that eliminating the pathology eliminates

the pain. Many, however, simply are not. Eliminating the pain is a distinct goal from eliminating the pathology.

In hopes of greater success, the dissatisfied have begun turning toward a mechanism-based classification of pain. This alternative approach, sometimes dubbed 'pain analysis,' seeks to classify types of pains (notice no longer pain *as such*) based on whatever underlying mechanism(s) determine(s) them. The cores assumption here is that successful treatment will result from manipulation of the activity of the mechanism(s) responsible for the pain. I argue that it is the difficulties for the development of pain analysis that supports that no mechanical model of any constructed type of pain is in the offing.

I have so far taken talk of pain "types" for granted, but it should now be asked: How should we type pain experiences? Etiology, pathology, history, and responsiveness to treatment are currently used though recognized to be insufficient by the scientific and medical communities for typing pains. The inadequacies of these traditional approaches are recognized even by those who think that the mechanism-based approach is too optimistic.[148]

The mechanism-based classification approach decrees that we should type pains by clusters of symptoms explicitly hoped to be correlated with different explanatory mechanisms. There is disagreement about the best categories, but all agree that different types of pain are going to have to be grouped separately if they are to correspond to underlying mechanisms. As seen in the previous section, neuropathic pain and inflammatory pain, for an obvious instance, are simply *not* determined or maintained by the same underlying mechanisms. The need to type pains separately by clusters of symptoms was recognized from the beginning of the movement. Recall from chapter 1 that Woolf, et al (1998), initially proposed 3 categories of pain—transient

[148] See, for example, Hansson, (2003).

pain, tissue injury pain, and nervous system injury pain—and *possibly* correspondent, explanatory, mechanisms. It was hoped that these correspondences would eventually be tightened-up so as to identify appropriate treatment targets. Since the 1998 editorial, other, similarly rudimentary, mechanism-based classifications of pains have followed.[149]

The centrality of correlating symptoms with mechanisms is depicted in the mechanism-based approach presented by Woolf and Mannion (1999):

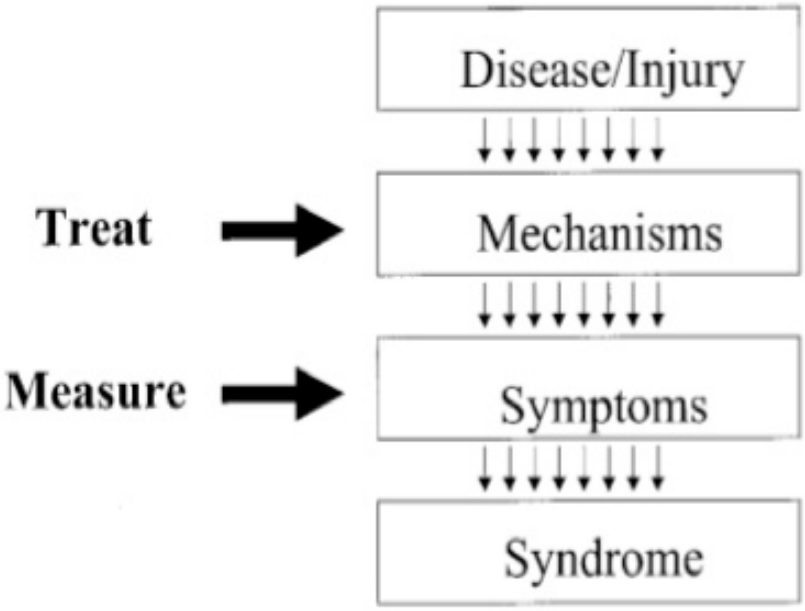

Figure 4.3: The Mechanism-Based Approach. Taken from Woolf and Mannion, 1999.

Though central, the biggest obstacle to the new approach is determining mechanisms on the basis of symptom clusters. As of now, there are no good correlations between types of pain and mechanisms—and the authors are explicit about this. The optimistic goal of mechanism-based

[149] Among others: Arner (1998), Gifford (1997), Jones (1995), and Lidbeck (2002).

classification is both unreached and necessary for the approach to get off the ground. As theoretically compelling as the new approach is, we cannot yet successfully implement it in practice.

In the landmark paper, Woolf et al. (1998), identify four criteria for evaluating the adequacy of a mechanism-based classification of pain: reliability, validity, generalizability and comprehensiveness. These criteria are referred to by almost everyone who has adopted the new approach. Reliability requires that the classification provides "truly operational criteria" and include appropriate "inclusion and exclusion criteria."[150] The basic idea is that the classifications need to illuminate targets whose manipulation will not manipulate too much or too little.[151] Validity is defined as ".. an estimate of the degree to which the classification system corresponds to the underlying biology of the disorder being studied."[152] Since there is no "gold standard" for this in pain research, we are encouraged to "… use an iterative, fallible process of searching for and identifying symptom clusters, biological markers, history, and treatment response."[153] If this sounds vague, that is because it is. The "gold standard" for pain is taken to be the subjective report, such that we can never know if we have correlated pain with the *biology* correctly, so we opt for quantity and correlate as many symptoms with as many mechanisms as possible. Generalizability requires that the classifications apply to the identified type of pain in all its forms—mild or severe, regardless of which disease it is presented as a symptom of. Comprehensiveness requires that there are no pains left out—a major problem with the

---

[150] Woolf et al., "Towards a mechanism-based classification of pain?," 228.
[151] Currently, "pain-killers" often affect much more than the mechanisms explaining any token pain experience; e.g., opioids have a widespread effect on functioning as there are opioid receptors throughout the entire nervous system. At the same time, other interventions, like antidepressants for chronic pain, affect both too much *and* too little. Mechanism-based classification is hoped to provide more tailored treatment targets.
[152] Woolf et al., "Towards a mechanism-based classification of pain?," 228.
[153] Ibid., 228.

monolithic models of pain discussed in the previous chapter. No one has provided a mechanism-based classification of pain that comes anywhere close to meeting these criteria.

To construct pain types following this approach, scientists and physicians must use the "symptoms and signs" presented by patients. They must, that is, use patients' reports. As Finnerup and Jensen (2006), point out: "Because patients reports symptoms rather than mechanisms of pain, the main clinical problem is to extrapolate mechanisms from symptoms or clusters of symptoms."[154] They identify three main obstacles to identifying mechanisms using patients' symptom reports. The centrality of these obstacles for my case against any type of pain being a natural kind makes the lengthy quote worthwhile:

> First, one symptom/sign can be caused by several initiating mechanisms. For example, cold allodynia [experiencing pain in response to non-noxious cold stimuli] might, speculatively, be caused by disproportionate loss of Aδ fibres and sensitization of cold receptors in peripheral neuropathic pain, whereas cold allodynia in central pain clearly has to be explained by other mechanisms.
>
> Second, the pain system is dynamic in nature. Neurobiological events change over time; for example, central sensitization [decreased thresholds of neurons] initiated by irritable nociceptors might eventually become independent of peripheral input. … This complexity is underscored by the fact that pain itself alters the processing of somatosensory information…
>
> Third, there is a temporal and anatomical sequence of events in the generation and maintenance of neuropathic pain, and a primary pathology might give rise to a series of secondary events… The symptoms can vary from day to day, perhaps because of plastic changes that depend on sensitization, inhibition and facilitation… If a treatment is effective in one of the final paths in this cascade of events—for example, acting on central sensitization or descending facilitation—it might be effective in several neuropathic pain conditions in which the symptoms and signs are caused by different primary mechanisms. The treatment might also modulate symptoms and signs without targeting mechanisms (110).

These problems are serious obstacles to developing a mechanism-based classification of pain—even, notice, of any constructed pain "types." So far, classifying pain into supposed "types" still

---

[154] Finnerup and Jensen, "Mechanisms of Disease: mechanism-based classification of neuropathic pain—a critical analysis," 110.

leaves double dissociations between symptoms and any mechanisms hoped to be explanatory of those symptoms. These problems are recognized by mechanism-based advocates and antagonists alike as being, so far, intractable.

Recall a summary of these problems from later in the same work, quoted first in chapter 1:

- To develop an effective mechanism-based classification, we need to be able to relate symptoms and signs to mechanisms and to identify specific treatments for specific mechanisms.
- One mechanism might give rise to different symptoms or signs; for example, upregulation of sodium channels in C-fibres increase fibre activity, resulting in burning pain, paroxysms and dynamic mechanical allodynia.
- Similarly, one symptom or sign can be caused by several initiating mechanisms; for example, cold allodynia can be attributed to different mechanisms in peripheral and central neuropathic pain.
- It is becoming clear that we have not yet obtained a viable mechanism-based classification for neuropathic pain, and more-rigorous studies are required to test this approach. ("Mechanisms of Disease: mechanism-based classification of neuropathic pain—a critical analysis, 113".)

Again: the double dissociations between symptoms and signs and underlying mechanisms, even *within* constructed types, is the largest problem facing the mechanism-based approach.

One core problem is that each token pain is the result of the activity of *numerous* mechanisms—genetic, chemical, biological, neural and social. Whether a particular individual experiences pain on a particular occasion is determined by the combined activity of these multiple mechanisms. This is now almost universally agreed upon by pain researches, and I am hard-pressed to choose a singular quote or reference supporting it. Here is just one, particularly clear, example from Smart et al (2008):

> The pathophysiological mechanisms of pain are not distinct entities but occur in a continuum. The generation, modulation and perception of pain involves simultaneous parallel processing in multiple and interrelated body systems, and as such, *any distinct categorization based on mechanisms is to some extent artificial since pain is the result of many mechanisms.* The extent to which varying combinations of mechanisms contribute to any patient's pain state can alter and

> change with time. There can also be great variation between individuals, injury and syndrome. (6, emphasis added)

The dilemma is that treatment is most effective when targeted at the mechanisms active in each particular case. For each token pain experience, however, there are *numerous* mechanisms that explain it—mechanisms whose interactions are dynamic (as opposed to static) over time.

Perhaps, however, we are trying to buy a cart when we should first secure our horses: Since moving from symptoms to mechanisms seems intractable, perhaps we should first identify mechanisms in isolation. If we could first identify mechanisms, we could then catalogue the symptoms that result from their workings—and identify them when presented by patients in the clinic. This approach, however, is even more hopeless than the last. Jensen and Baron (2003) lucidly discuss the problem:

> Most likely symptoms, signs and mechanisms are related, but probably not directly. For example, a diabetic patient may have: steady pain, touch evoked pain, paroxysms [sudden onset of symptoms] and non-painful paraesthesia [sensation of e.g., tingling or numbness]. In these cases, several mechanisms can be involved such as tissue injury due to ischaemia, sensitization of peripheral receptors, ectopic activity in sprouting regenerating fibres, phenotypic changes in dorsal root ganglion (DRG) cells, spinal reorganization etc. In other pain states, one mechanism may give rise to different symptoms and signs. For example, a peripheral nerve entrapment may cause paroxysms due to ectopic activity from the lesioned nerve and an associated (extra)-territorial brush evoked pain due to C-fibre evoked central sensitization. Thus symptoms and mechanisms involved in a particular pain condition cannot always be predicted. Is it then impossible to dissect pathophysiological mechanisms? It is probably not possible. As patients report symptoms and 'not' mechanisms and assessors disclose signs and 'not' mechanisms, the starting point should be the patient with a focus on how symptoms and signs translate into mechanisms rather than vice versa. ("Translation of symptoms and signs into mechanisms in neuropathic pain," 2-3.)

Numerous examples are offered here to support the claim that moving from mechanisms to symptoms is not promising ("probably not possible") and has no known advantage over moving from symptoms to mechanisms. The problem, here again, is that each token pain experience is explained by the differently weighted activation of *many* mechanisms. Knowledge of the activity

of any single, identified mechanism will not yield successful prediction of pain symptoms; the products of the mechanistic activity will vary depending on the other mechanisms involved in each particular case.

Until we can arrive at better correlations between clusters of symptoms and underlying mechanisms, the new model is no better than the old medical model for reliability, validity, generalizability or comprehensiveness. I submit that this new approach will not do better because the mechanistic activity determining any given token pain cannot be reliably generalized about to explain *other* token pains.

The variability of individuals' token pain experiences given the same diseases, same stimulation, *same mechanisms*, and same everything-we-can-think-of-to-control-for is, unsurprisingly, a serious problem for effective treatment. Since treatment is the goal of pain research, the fact that it is *individuals* who are treated cannot be overemphasized. As Woolf and various collaborators have pointed out numerous times (1998, 1999, 2001) our inability to adequately deal with the highly idiosyncratic nature of token pain experiences is exacerbated by the fact that pain research almost always glosses over individual differences in, for instance, drug trials and neuroimaging. This is problematic for pain treatment since there is extensive individual variability, not only in laboratory trials, but much more so in the clinic. Nonetheless, it is the individual that must be treated.

But consider: What more could one ask from scientific research? Insofar as science provides generalizations, it deals with types, not tokens. Research necessarily glosses over individual differences to generalize about treatments, brain activity, or anything else under investigation. Glossing over of these individual differences, however, blocks reliable correlation between mechanisms (discovered by the researchers) and pain symptoms (presented by the

individuals). *Adequate explanation of a token pain experience requires an individualized explanation.* Individualized explanations include the genetic, molecular, chemical, biological and social factors relevant to a *particular* individual's pain at a *particular* time. Explaining any token pain experience crucially relies on these individual differences; adequate explanations of pain are unavoidably idiosyncratic.

It is, finally, for this reason that no type of pain is itself a natural kind: the unavoidable idiosyncrasy of adequate explanation blocks successful generalizations. No known implicated mechanism is explicable of any type of pain, and no mechanism yet to be identified is going to be. Why? Because each token pain experience is the result of the idiosyncratic combination of activity across many implicated mechanisms. More mechanisms implicated in pain experiences are being discovered all the time, and provide novel treatment targets. That is useful and important research. The role that the identified mechanisms play in any individual case, however, is also determined by activity in numerous other causally relevant mechanisms—this combination is idiosyncratic.

In sum then, there are two reasons no pain "types" are not natural kinds. First, no one mechanism explains a token pain experience; each token pain experience is explained by the activity of multiple mechanisms. This by itself may show that no pain "type" is a natural kind; token pain experiences are explained by multiple mechanisms such that no method of typing the tokens will yield a type whose properties are explained by a single, underlying mechanism. Second, the activity and interaction of those multiple mechanisms is idiosyncratic. Science's generalizations will, inevitably, abstract away from the individual differences—but those differences are crucial in the explanation of each token pain experience. Generalizations about pain "types" are not useful; in particular, they are not adequate for treatment.

To be clear: I expect that we will find more and more mechanisms implicated in pain experiences that we can successfully generalize about. We will undoubtedly come to better understand, for instance, how the different systems of the body interact, how different receptors are facilitated or inhibited by various proteins, and much more.

To be emphatic: None of this mechanistic knowledge will yield an identification of pain "types" that are usefully referred to in science's generalizations. Each token pain is the combined result of an individual's particular genetics, chemistry, biology, memories, upbringing, culture, and much else that we are still trying to understand. No single, underlying mechanism explains all of that. Explanations of token pain experiences are not reliably generalized about, because pain is determined by an *idiosyncratic* convergence of activity across multiple mechanisms.

Though the intractability of the obstacle it presents for scientific generalizations about pain is not, the individuality of pain experiences is recognized by leaders in the field. Recall from chapter 3 that Melzack and Wall referred to each pain experiences as *unique*, Price took the meaning each *individual* attaches to their experience to be crucial to pain's affective components, and the neuromatrix theory of pain posits an *individualized* neurosignature. Attention to the individuality of pain also illuminates what is right about those definitions of pain, like the influential one offered by the IASP, that define pain *subjectively.* Such definitions rightly place a particular subject and their particular experience at the center, and undermine predictions about who will feel pain when.

One might object that pain experiences only require an idiosyncratic explanation in the "odd" cases. Surely "normal" pain is not idiosyncratic, even if other strange cases are. Almost everyone will feel pain if they get chili powder in their eyes (chemical), put their hand on the stove (thermal), or get hit with a hammer (mechanical). The best candidate for "normal" pain is

acute pain experienced as the result of brief noxious stimulation presented in an antecedently painless state: nociceptive, or, physiological pain. Here too, however, individuality looms. Nociceptive pain is subject to a \host of cognitive and environmental factors.[155] That these factors are causally relevant to (if not actual constituents of) nociceptive pain is well known. As emphasized in the previous chapter, all contemporary scientific models of pain recognize a cognitive component. That at least this component is idiosyncratic is obvious.

The cognitive dimension of pain experience fosters a strong intuitive case for the idiosyncrasy of token pain experiences. Unlike pain scientists, however, philosophers tend to take this cognitive dimension as a *reaction* to pain. From this point of view, the idiosyncrasy of the cognitive dimension is weak support for the idiosyncrasy of the pain experience—supporting only an idiosyncratic *reaction* to pain. That the cognitive dimension of pain is causally relevant to the sensory experience, however, is hard to deny for anyone with the slightest familiarity with the pain literature. Beliefs, memories, and expectations sometimes alter sensory experiences such that a noxious stimulation is not experienced as pain, or is experienced as much *more* painful than it would be with a different set of beliefs, memories and expectations. Explaining or treating pain often requires manipulating the cognitive mechanisms that partially determine it.

For those unconvinced, however, the role of genetics in nociceptive pain experiences may be more persuasive. The search for pain genes provides a nice case study—of the explanatory role of multiple mechanisms even in "normal" pain and the idiosyncratic activations and interactions of those mechanisms. Genetic research suggests that sensitivities to particular types of noxious stimulations are reliably correlated with genetic differences—at least in the mouse (Mogil et al., 1999a, 1999b, 2000, Lariviere 2002). These genetic differences correlate not just

[155] See e.g., Pennebaker, 1982, for a number of fascinating experiments illuminating this.

with nociceptive pain, but five (and growing) identified pain subtypes. Do these genetic differences provide a mechanistic explanation for pain? For nociceptive pain? Not hardly.

Genetic differences alone will not explain a person's pain experience—for that, we must take account of the activity in the other mechanisms implicated in "normal" pain experiences and their interactions with the newly identified genetic mechanisms. The authors are explicit that identifying pain genes is the job of identifying genes *relevant* to pain. The role of genetic mechanisms in explaining pain is always discussed in the modest terms of modulation. Here is one choice quote:

> A number of benefits are likely to derive from the study of pain genetics. First, the discovery of novel genes and the confirmation of known genes as being pain-relevant should spur targeted drug-discovery efforts. Second, the elucidation of genetic polymorphisms rendering individuals with varied sensitivity to pain and its inhibition should allow idiosyncratic use of existing therapies, maximizing their effectiveness and minimizing side-effect liability. … It should be emphasized, however, that environmental factors, acting both alone and in concert with genotype, exert more overall influence on complex biological phenomena, such as pain, than do genes… thus, in addition to putting up with the pain required to decipher pain genes, the search for sociocultural and other nongenetic factors that influence pain susceptibility and response must be pursued with equal vigor. ("Pain Genes?: Natural Variation and Transgenic Mutants," 803.)

Identifying genetic mechanisms implicated in pain experiences is important and may provide avenues for novel treatments, but it will not provide a mechanical model for pain or any pain "type"—including "normal" pain. In the above passage, sociocultural and environmental factors are claimed to be even more relevant. Numerous other mechanisms have, of course, been identified throughout this chapter and the last, and there are more that I have lacked the space to discuss. No *one* of these explain any token pain, much less any "*type"* of pain artificially constructed out of those tokens. Genetic mechanisms, as the authors conclude, appear to be "pain-relevant." They are not, however, pain-explanatory. Again: Each pain experience is explained by the idiosyncratic activity of multiple mechanisms in a particular time and place.

The claim that each token pain experience is idiosyncratic should not be overblown. Many people have many things in common genetically, chemically, biologically, and socially. For most people, most of the time, aspirin will help a headache. Why? Because aspirin, as above, targets inflammatory responses. Most people have enough in common biologically, genetically, chemically and socially that their idiosyncratic combination of relevant mechanisms will *include* activation of inflammatory responses when they have a headache. This mechanism is targeted by aspirin so, in some cases, voilà, no more pain. Do we then have a mechanistic explanation of a pain type, headache? Not on your life. Headaches continue to be one of the most befuddling pain "types." We know inflammatory mechanisms are implicated, but we also know that inflammatory mechanisms are not the whole story. That is why sometimes, when you have a headache, aspirin *does not* help. For some people, moreover, aspirin *never* helps and for most people a placebo will help just as well. Why? Because *other* mechanisms are also involved every time you have a headache—sometimes they are more determinate than others. Sometimes the pain simply will not abate until you, for instance, reconcile with that disgruntled coworker or sincerely *expect* that it will abate. There is nothing spooky here. There *are* numerous mechanisms, including cognitive mechanisms that determine whether you feel pain—even common headache pain. That is why treatment must be individually tailored.

Pain research, I submit, will better advance by focusing on particular mechanisms *implicated* in pain experiences, and less on figuring out which *pain types* are correlated with which mechanisms. Why? Whatever cluster of symptoms we decide to deem a pain type, that "type" will not correlate well with the activities of any particular mechanism, and the convergence of the activity of multiple mechanisms explaining one token pain will not reliably generalize to others. That is what mechanism-based classification advocates are finding.

It may be objected that whether any pain "type" has a mechanistic explanation is best settled by the scientists. Science should not be done from the armchair, and I admit that I am a bit uncomfortable making the claim. The support for it, however, comes from the science and not simply from the stigmatized chair. It is now well known (among the scientists if not the philosophers) that each pain experience results from the activity of multiple mechanisms. It is also well known that the convergent activation of these mechanisms is idiosyncratic. My contribution is meant only to make explicit that this idiosyncrasy of multiple converging mechanisms blocks useful generalizations about pain or any pain "types." It is the fruitful identification by the scientists of so many mechanisms and their idiosyncratic interaction that allows the non-naturalness of pain and pain "types" to be revealed.

Advocates of mechanism-based classifications are, of course, still hopeful that they will identify distinct clusters of symptoms that reliably correspond to distinct mechanistic activity. They are optimistic that we will eventually arrive at useful generalizations about pain "types"—though they have given up on pain as such. The heart of the problem with this optimism is that individual differences are central to explanations of pains, but individual differences get lost in scientific generalizations. Science delivers generalizations, but token pain experiences, and their mechanistic explanations, are not reliably, comprehensively, or validly generalized. Pain tokens aren't usefully typed for scientific inquiry.

In closing this section, I would like to make plausible that the view that each token pain experience is idiosyncratic is a natural next step in the development of pain science. Recall the brief history presented in chapter 3: pain science progressed from a monolithic view of pain (still holding sway in philosophy), centered on fiercely debating which unitary phenomenon pain is, to convergent views that characterize pain as a multidimensional experience, to attempted

categorizations of pains into types by pathology, and finally, by mechanism. We first had specifity theories and theorized that pain was a distinct phenomenon subserved by specialized systems—we believed in pain receptors, pain pathways and pain cortical centers. Gate-control theory disillusioned us of this simplicity by discrediting the idea of receptors, pathways, or cortical centers specific for pain. In the wake of freedom from specificity theory, numerous mechanisms never before thought relevant to pain were confirmed to be causally relevant and susceptible to pain-relieving treatment. More mechanisms than Melzack and Wall had initially dreamed. We know now that different mechanisms are implicated in different pain experiences and that multiple mechanisms are implicated in each one. We have most recently learned that these mechanisms are susceptible to idiosyncratic factors—genetics, memories, past trauma, cognition, and more. This is the history of pain science. It is time for the next step in conceptualizing the nature of pain: Each pain experience is idiosyncratic and has an idiosyncratic explanation. Not only is pain a convergent phenomenon—it is an *idiosyncratic*, convergent phenomenon. This is the direction pain science has been moving and it is time to admit it, and its implications, outright. One can see the movement toward an understanding of pain as idiosyncratic in the rise of the modern pain clinic—where individualized, multi-dimensional treatment is at the core.

Most important for present purposes, the idiosyncratic convergence of multiple mechanisms is why there is no mechanistic explanation of pain or any pain "type." Many mechanisms implicated in pain experiences will undoubtedly continue to be discovered. As *evidenced* by these discoveries, however, no pain or pain "type" will submit to a mechanistic explanation. That is why generalizations about pain have inevitably been, and will continue to be, unreliable. It is also why pain treatment is so woefully inadequate. As Melzack and Wall said

some time ago, each token pain experience is unique. Neither pain nor any "type" of pain is a natural kind.

4. Objection

An important objection may have been long brewing. In chapter 1, I presented the scientific utility criterion, which stated that natural kinds are those kinds that are usefully referred to in science's generalizations. References to those kinds will be more or less useful, and the kind referred to thereby more or less natural, I argued in accordance with the success of the explanations and predictions facilitated by the references. Following Quine and Boyd, I then claimed that such kinds will, at least in principle, be mechanistically explicable. That a kind be useful for science *and* that it submit to a mechanistic explanation were obviously intended to be harmonious criteria for natural kindhood; i.e., though some kinds admitting of mechanistic explanations may not be usefully referred to in any of science's generalizations, there should be no kind that is usefully referred to in the generalizations of the relevant science that does *not* admit of a mechanistic explanation in principle.

The pain research I have canvassed may suggest that pain is a counterexample to the harmony of these two natural kind criteria. I concluded that there is no mechanistic explanation of pain or pain "types." Nonetheless, the term 'pain' has been instrumental in the development of pain science, and that science has borne many fruits. Denying either of these claims would be absurd. Investigating pain has increased our understanding of the brain, the nervous system, primary afferents, homeostasis, immune responses, and much more. Perhaps most importantly, these investigations have resulted in numerous novel treatments.

This objection is important, but it is importantly misplaced. It has been very useful to investigate pain, but this does not show that the term 'pain' is usefully referred to in the generalizations of any relevant science. Before elaborating this point as applied to pain, I consider two other terms from the history of science: 'aether' (i.e., lumiferous aether) and 'humors.' Investigation into aether and the humors was fruitful, but their referring terms are *not* usefully employed in science's generalizations.

First, briefly consider lumiferous aether or simply *aether*. 'Aether' was originally used to refer to a singular substance, made up of numerous particles that facilitated the travel of light. Many investigations of aether were carried out—investigations that bore fruit! Over time, however, science began to attribute increasingly complex and seemingly conflicting properties to aether, and it became hard to see how anything could *have* all of those properties. It was Einstein who seems to have eventually refuted its existence. It is interesting to note that he pointed out that the scientific community could, if they wished, redefine 'aether' to capture something similar but new. The scientific community deferred and the term ceased to be employed in scientific inquiry.

Let us imagine that the scientific community was correct in dropping the term and explain it by the disutility of 'aether' for further inquiry. We would not thereby be denying that the *investigation* of aether was useful for the advancement of science. It is consistent with this that the term 'aether' was rightly dropped because its employment in scientific generalizations was not useful. According to the scientific utility criterion, insofar as employment of that term is *not* useful when so employed, aether is not a natural kind. 'Aether' may have been, for some time, the best the science could do. It was the useful investigations of aether itself, however, that

eventually showed that the term ‘aether’ did not refer to anything useful for scientific inquiry. Physics could do better.

Now consider ‘humor’. On into the 20th century it was widely believed that there were four basic substances in the body called the humors: black bile, yellow bile, phlegm and blood. Health was explained by the their balance and their imbalance was theorized to result in various maladies. This theory, called humorism, is ancient and litters the history of medicine. It is, of course, now accepted that humorism is false and that ‘humor’ has no referent useful for scientific inquiry. As with Einstein’s suggestion about ‘aether,’ notice that we could, if we wanted, employ the term ‘humor’ to refer to something similar but distinct—blood, perhaps. The scientific community has again, however, opted to drop the term altogether. Enough of humorism has been discredited that the term ‘humor’ is presumed useless for scientific inquiry.

Notice again that many investigation of the humors were carried out—investigations that bore fruit! We should not deny that the *investigations* of the humors were useful. This is perfectly consistent with the scientific utility criterion, which states that insofar as ‘humor’ is *not* usefully there employed in scientific generalizations, its referent is not a natural kind. Reference to the humors may have been, for some time, the best that science could do. Useful investigations of the humors, however, eventually showed that the term ‘humor’ did not refer to anything useful for scientific inquiry. Medicine and biology could do better.

Here we should be reminded, as discussed in chapter 1, that different terms may be more or less useful at different times. Science sometimes progresses from an understanding of a kind as a Russell, to a Mill, to a Pierce kind. When this happens, our understanding of a kind alters and references to that kind become *more and more* fruitful for induction and explanation. We continually refine comparative similarity spaces in the light of subsequent inquiry. For this

reason, following Quine, I claimed that a kind's naturalness is a matter of degree. In the case of the humours and luminferous aether, however, references to those kinds became *less and less* fruitful for induction and explanation. Scientific investigation of those kinds eventually revealed that the relevant sciences would do better to cease referring to them altogether.

The varying utility of different terms for different inquiries at different times forces a decision concerning the naturalness of the referents of a term once we have come to believe that the term is no longer usefully employed in scientific generalizations. Ought we to say that natural kinds are those that ever were or ever will be usefully referred to or only those usefully referred to at the time of decision? I have chosen to develop the theory of natural kinds in chapter 1 in accordance with the latter decision. Including kinds as natural only for as long, and insofar, as they are usefully referred to in science's generalizations has benefits. In particular, it allows me to say, as is intuitive, that we have been wrong, and may currently be wrong, about which kinds are natural. We should stand behind our best theories at, but only at, any given time.

Consider again aether and the humors. Investigation into these was useful. Employing the scientific utility criterion, it may be tempting to say that these kinds are natural to recognize the seeming contribution of reference to them. Doing so, however, conceals that reference to these kinds ultimately became less and less fruitful in the light of inquiry—eventually becoming useless. Rather than say that reference to the purported kinds was useful, it is better to say that the relevant investigations were useful *despite* these references. Given what we now know, we should not credit reference to those kinds with the utility of the investigations; credit belongs to the things the relevant investigations allowed us to discover despite those references.

Investigation of a natural kind allows for better and better explanations and about that kind. When investigation, even fruitful investigation, reveals that a kind is not usefully referred

to in science's generalization, that kind is not natural. We might have thought it was, but we were wrong. Such was the case with aether and the humors. Discovering this allows us to say that aether and the humors were *not* natural kinds after all. We just, errantly, thought they were.

Returning to pain: The investigation of pain has been very fruitful. I do not deny the utility of those investigations. As with the humors and the aether, however, that does not establish that pain or pain "types" are natural kinds. We may, if we like, say that reference to pain *was* useful for science. I should rather say that these investigations were fruitful *despite* their references to pain, just as I should like to say of aether and the humors. Whichever we say, however, it is time to recognize that the useful investigations of pain has itself showed that neither pain nor pain "types" are natural kinds. If reference to these was once useful for scientific inquiry, it is useful no longer. Biology, anatomy, physiology, pharmacology, anesthesiology (and many other relevant sciences) can all do better.

Despite its non-naturalness, however, pain is a real and important phenomenon in our daily lives. In this way pain is *different* from aether or the humours. Accepting the existence of pain and the utility of reference to it for everyday life may suggest a pluralism about pain: though pain is mechanistically heterogeneous, it is still usefully referred to in scientific generalizations either because it is homogeneous at another level of description or because the mechanical differences are negligibly related to the success of the generalizations. I argue in the next chapter that pluralism should be resisted. The idiosyncrasy of pain experiences blocks utile reference to pain or pain "types" in the generalizations of any relevant science. Avoiding pluralism, however, may suggest accepting eliminativism: if pain is not usefully referred to in science's generalizations, then it doesn't exist. In the concluding Chapter 6, I argue that this too should be resisted. While science ought to drop the term 'pain', the layman ought not. Pains exist

in any interesting sense of ‘exist,’ and referring to them in everyday life is sometimes useful. The final two chapters, then, argue against both pain pluralism and eliminativism despite what I now take to be established: Neither pain nor any “type” of pain is a natural kind.

Chapter 5: Pluralism

1. Introduction

I argued in the previous chapters that neither pain, nor any pain "type" are natural kinds, because they are not explained by a single, underlying homeostatic mechanism. The lack of an underlying mechanism leaves references to pain or pain "types" in science's generalizations lacking in utility. Furthermore, the idiosyncrasy of token pains undermines the prospects for future discovery of an as yet unknown mechanistic explanation of pain or pain "types." I call the view that neither pain nor any "type" of pain is a natural kind the non-natural kind thesis. This chapter distinguishes the non-natural kind thesis from pluralism about pain and argues for its superiority.

I define pluralism about a natural kind the claim that despite mechanistic differences across the kind's members, they nonetheless constitute a natural kind since they are nonetheless usefully referred to in (at least some) generalizations of (at least some) relevant sciences. This may happen for one of two reasons: either the kind is homogenous at one (usually "higher") level of description though mechanistically heterogeneous at another (usually "lower" level), or the identified mechanical multiplicity is negligible for the success of generalizations about it. These reasons are the same when a kind's mechanical heterogeneity at one level is negligible relative to the homogeneity at another.

Though I established that there are mechanical differences among token pain experiences, even among purported pain "types," pluralism about pain is not thereby refuted. If the identified mechanical differences are negligible relative to *some* generalizations, or if there is some level of description at which there is sufficient mechanical homogeneity, then the non-natural kind thesis should be rejected in favor of pluralism. It is admittedly unintuitive that something as central to our lives as pain is not a natural kind. Insofar as pluralism allows for *some* mechanistic

differences across tokens of natural-kind types, it may be compatible with the arguments presented in the previous chapters, without succumbing to the unintuitive non-natural kind thesis. Whether we should be pluralists about pain turns, in sum, on whether is there *enough* homogeneity among pains or pain "types" to underwrite *some* useful generalizations *despite* the identified mechanical differences. How can we tell?

Help can be garnered from the recent debate between Weiskopf (2009) and Machery (2009) concerning pluralism about concepts. Weiskopf and Machery both accept a picture of natural kinds informed by Boyd and nearly identical to the scientific utility criterion developed in chapter 1. They are, furthermore, in agreement that there are no properties common to all concepts, and that different mechanisms may determine different token instances of concept deployment. Nonetheless, they disagree about whether the evidence supports that concepts form a class that is usefully referred to in psychology's generalizations. They disagree, that is, about whether concepts are a natural kind.[156]

This chapter employs lessons learned from this debate to argue that pluralism about pain and pain "types" should both be resisted. In section 2, accordingly, I look at Machery's argument for the non-naturalness of concepts, and compare it to my own for the non-naturalness of pain or pain "types." In section 3, I briefly consider Weiskopf's objections to Machery and elucidate their relevance for evaluating the non-natural kind thesis. In section 4, I turn directly to the question of whether we ought to reject the non-natural kind thesis in favor of a pain pluralism. I argue that the mechanical differences between token pains established in the previous chapters are *not* negligible for the utility of the relevant generalizations, and there is no level of

[156] I consider the debate in some detail, but my intention is to remain neutral about whether concepts are a natural kind.

description at which reference to pain or pain "types" *is* useful for scientific generalizations. Hence, whatever the case for concepts, pluralism about pain is misguided.

2. Machery

Machery (2009) develops a position concerning concepts that he calls "scientific eliminativism." According to this view, 'concept' should be eliminated from scientific discourse due to the heterogeneity between what he takes to be fundamentally different types of concepts. The position on pain and pain "types" developed in the previous chapters bears a striking resemblance to the position on concepts developed by Machery.[157] Though remaining neutral about Machery's arguments, the relevant considerations for whether we should accept pluralism about concepts are also relevant for determining whether we should accept pluralist about pain or pain "types."

Taking psychology to be the relevant science whose generalizations refer to concepts, Machery presents an argument for eliminating the term 'concept' from psychological inquiry summarized in the following five points:[158]

> 1) The best available evidence suggests that an individual typically has several concepts for each category of objects.
>
> 2) Coreferential concepts have very few properties in common; i.e., coreferential concepts belong to very heterogeneous kinds of concept.
>
> 3) Evidence strongly suggests that prototypes, exemplars, and theories are among the fundamental kinds of concepts.
>
> 4) Prototypes, exemplars, and theories are often used in distinct cognitive processes.

---

[157] Notice that Machery is a scientific eliminativist about concepts but not concept types, whereas I advocate a scientific eliminativism about pain *and* pain "types."
[158] Adapted from Machery, *Doing Without Concepts*, 52-53.

> 5) The notion of concept ought to be eliminated from the theoretical vocabulary of psychology because it might prevent psychologists from correctly characterizing the nature of our knowledge in long-term memory and its use in cognitive processes.

Notice, as Machery himself is aware, that this argument is straightforwardly empirical as against (merely) conceptual. 1-4 are all purported results from psychological, scientific, inquiry. As such, future inquiry may overturn this evidence and therefore the eliminativist claim embodied in 5 that rests on it. The empirical structure of this argument is appropriate since, as I argued in the previous chapter, the claim that a term should be eliminated from scientific inquiry depends on evidence *from* scientific inquiry.

Machery's criterion for naturalness is very similar to that developed here in chapter 1. Machery claims that "… the notion of a natural kind singles out those classes about which scientifically relevant inductive generalizations can be formulated."[159] So too, Machery also turns to the homeostatic mechanisms emphasized by Boyd for indicators of which classes are likely to support inductive generalizations.

Nevertheless, there is a slight but important difference between Machery's account and my own. Machery holds that a non-natural kind may nonetheless sometimes be useful for scientific inquiry:

> A theoretical term that has been found to fail to pick out a natural kind should be kept if it plays a useful role. Not all terms in science are assumed to pick out natural kinds. Terms that do not pick out natural kinds might have several functions… However, by the same token, a theoretical term that has been found to fail to pick out a natural kind should be eliminated if it fails to play a useful role or if it plays a harmful role. I believe that the latter is likely to be the most common case. When evidence emerges that a hypothesized natural kind term fails to pick out a natural kind, keeping this theoretical term is likely to prevent the development of a new classification system that would identify the relevant natural kinds (*Doing Without Concepts*, 239).

---

[159] Ibid., 231.

It is a bit puzzling why Machery accepts that some non-natural kind terms may nonetheless be useful *for scientific inquiry*. The offered suggestion is that such a term my be a useful shorthand; e.g., 'concept' may be a useful shorthand for referring to bodies of knowledge used by default in the process underlying most higher cognitive concepts. Machery holds that concepts *not only* do not form a natural kind, but that this shorthand is not useful.

If a term *is* a useful shorthand, however, it is unclear why the term and its referent are not thereby natural. Naturalness, for both Machery and me, is determined by a kind's role in facilitating generalizations. Non-natural kind terms that Machery thinks should nonetheless be retained would have to be useful shorthands while nonetheless *not* usefully supporting any generalizations. If the shorthand is not useful by way of facilitating generalizations, what is it useful to scientific inquiry *for*? On this, Machery is mute.

Though utility is extrinsic to Machery's criteria for naturalness, it is nonetheless a crucial step in his argument for scientific eliminativism. As I take utility of reference to a kind in scientific generalizations to be the core determinate of its naturalness, Machery's argument for scientific eliminativism makes into two steps what I have made only one. Machery establishes the non-naturalness of concepts and *then* argues further that retaining the term 'concept' is useless and perhaps harmful, whereas I take the lack of utility in referring to the kind and the non-naturalness of the kind to be established by the same evidence. As both Machery and I rely on Boyd's criteria to indicate natural kinds, however, and since Machery thinks concepts are thereby non-natural *and* the term 'concept' not usefully retained, I put this difference in our approaches aside in what remains.

Of utmost importance for present purposes, are Machery's clear distinctions between scientific eliminativism, pluralism and traditional eliminativism. Pluralism about a kind is the

claim that it is natural despite mechanistic differences between its members. Scientific eliminativist arguments conclude that a kind is non-natural, and retaining the term used to refer to it is not useful for any scientific purposes.[160] Traditional eliminativist arguments go further by concluding that the relevant term never successfully refers and should be eliminated from everyday discourse, whereas scientific eliminativism is neutral concerning these additional claims.

Machery construes traditional eliminativist arguments as having the following form:

1) X is Y by definition.
2) There are no Ys.
3) Therefore, there are no Xs.

He takes issue with the *form* of these arguments. The problem, Machery claims, is that their success relies on which theory of reference is correct, and "… eliminativism should not be hostage to debates about reference."[161] To establish that claim he considers the relevance of both causal-historical theories and description theories of reference.

Machery seems to think that if causal-historical theories of reference are adequate, then arguments of the traditional eliminativist type will *never* be successful. Reference to the kind, according to Machery, would then depend on what is historically associated with the kind term and subsequent evidence will never overturn this. Taking belief as an example, he says: "If 'belief' was historically associated with a class of mental states, 'belief' refers to this class, whether or not the generalizations that define 'belief' are true."[162] This is too hasty. Causal-historical accounts require that employments of the relevant term bear the *right sort* of causal

[160] Again, according to my scientific utility criterion this second conjunct is a function of the first.
[161] Ibid., 230.
[162] Ibid., 225.

relations to the term's referents—with "right sort" being a matter of well-known matter of dispute. Returning to Machery's example, if tokenings of the concept of BELIEF are not caused in the right sort of way by *beliefs*, then 'belief' does not successfully refer to beliefs when uttered—at any rate, so says the causal-historical reference theorist. Crucially missing from Machery's discussion here is the *causal* part of the causal-historical account. If Machery intends the "historical associations" he discusses to *include* the crucial casual part of these theories, then his construal of the causal-historical account begs question in the current context. On his construal, that is, the causal-historical account denies the very possibility of premise 1 and 2 of the traditional eliminativist arguments based on an antecedent denial of the conclusion. Surely this is not true for all X, and no sane causal-historical reference theorist should accept that it is.[163]

Machery also seems to think that if a descriptive theory of reference is correct, then whether a traditional eliminativist argument is successful depends *solely* on whether anything satisfies the description. This too is too hasty. Left out of Machery's discussion is the option to revise the description. Though beginning with the belief that X is Y by definition, if it is discovered that there are no Ys, the descriptivist is not forced to conclude that there are no Xs. They may instead *revise* their description of Xs, thereby rejecting the first premise that Xs are Ys by definition. That is, the descriptive theorist may always revise their description of Xs such that something satisfies the referent. If reference to X would be useful subsequent to revision,

[163] As an example, consider a traditional eliminativist argument for the non-existence of angels in Machery's identified form: 1) All angels are non-physical entities by definition. 2) There are no non-physical entities. Therefore, 3) There are no angels. As causal-historical theories of reference require that successful reference of 'angel' depends on the right causal connection between tokenings of the concept ANGEL and angels, then the truth of 1 and 2 are reasons for denying that 'angel' ever successfully referred when uttered. This conclusion may be happily accepted by the even the staunchest causal-historical theorist.

moreover, it seems to me that this is exactly what the descriptivist *should* do.

Despite these problems in Machery's discussion of theories of reference, he is quite right about the form of traditional eliminativist arguments. Most important, he is quite right to distinguish them from what he calls *scientific* eliminativist arguments.

Machery formulates two kinds of scientific eliminativist arguments—vertical and horizontal. Both conclude that a term ought to be eliminated from scientific discourse. Vertical arguments purport to show that the target term picks out *multiple* natural kind, whereas horizontal arguments conclude that it picks out *no* natural kinds.

Vertical arguments, accordingly, have the following form:

> 1) The theoretical term in question features in generalizations that involve different types of properties.
>
> 2) These distinct types of generalization do not line up with each other. (i.e., they have different extensions/can be applied to different kinds of organisms, yield conflicting predictions/generalizations, etc.)
>
> 3) Therefore, eliminate (from scientific discourse).

Horizontal arguments for a theoretical term 'K,' by contrast, have the following form:

> 'K' does not pick out a natural kind to the extent that A,B, and/or C are satisfied:
>
> A) There are very few generalizations true of Ks beyond the identifying properties of Ks and there *are* many generalizations true of the members of *subclasses* of K.
>
> B) Generalizations thought true of Ks are in fact true of members of a *superset* of Ks and are true of K because Ks belong to this superset.
>
> C) The generalizations assumed to be causally grounded are accidental.
>
> Therefore, eliminate (from scientific discourse).[164]

[164] Notice that the recommendation for term-elimination proffered by both types of scientific eliminativist arguments is based on consideration of the success of the generalizations, instead of determining naturalness on these grounds and determining the term's retention separate. This is,

Machery's own argument is a horizontal scientific eliminativist argument. A, B and C are all argued to be true to a great extent when the variable 'K' is satisfied by concepts. 'Concept' should therefore be eliminated from scientific vocabulary. Notice that it does *not* follow that 'concept' never successfully refers when uttered, or that 'concept' should be eliminated from our everyday discourse.

Having distinguished his scientific eliminativism about concepts from traditional eliminativism, Machery is then careful to distinguish it from pluralism. Concept pluralism is attractive if, despite many of the generalizations about concepts being underwritten by different mechanisms, they nonetheless have some properties in common that underwrite yet *other* generalizations unaffected by the mechanistic differences. Machery is particularly vulnerable to a pluralist appeal since he accepts that prototypes, theories, and exemplars are fundamental *kinds* of concepts that are themselves natural kinds. This suggests that concepts are themselves a natural kind, with other natural kinds recognized by Machery nested within it: namely, prototypes, theories, and exemplars. Once these different types of concepts have been identified, how do we decide whether to go scientific eliminativist or pluralist about concepts proper?

Machery argues that whether concepts proper are a natural kind *in addition* to prototypes, theories, and exemplars depends on whether there are useful generalizations about concepts that we are unable to make without 'concept' in our scientific vocabulary. Does 'concept' underwrite any generalization in psychology that 'prototypes,' 'exemplars,' and 'theories' could not facilitate better or as well? Machery argues in the negative. He urges that the differences between the fundamental types of concepts undermine the success of generalizations about concepts

---

I think, as it should be, but is in tension with Machery's earlier methodological claims that a term may be usefully retained in scientific discourse despite its referent's inability to successfully support inductive generalizations.

proper, such that concepts are not appropriately referred to as a unified class in any scientific generalizations—no matter the level of description. Prototypes, theories and exemplars are therefore natural kinds and concepts proper are not.

Notice that the issue is *not* whether the types of concepts Machery accepts as natural have *anything* in common. Of course they do:

> Clearly, prototypes, exemplars, and theories have plenty of irrelevant properties in common, for example, *not being cups of tea*. What matters, however, are those properties that psychologists take to be important when they characterize concepts. As we saw repeatedly, psychologists focus primarily on what knowledge is stored in concepts and on how concepts are used in cognitive processes. If exemplars, prototypes, and theories store different types of knowledge that are used in different cognitive processes, this would suggest that they have little in common. (*Doing Without Concepts*, 118.)

Which properties are scientifically relevant is determined by what psychologists are trying to explain and predict. What, that is, the relevant scientists are attempting to generalize *about.* Machery thinks that psychologists expect that concepts will indeed have enough such properties in common to underwrite *some* successful generalizations:

> … psychologists typically expect that over and beyond the differences between concepts and between kinds of concept, concepts (or, maybe, most concepts) share many scientifically relevant properties. In other words, they expect the class of concepts to yield numerous inductive generalizations—principally about what type of knowledge is stored in concepts and how they are used in cognitive processes, but also about what their format is, how they are acquired, and where they are localized in the brain… (*Doing Without Concepts*, 54.)

Machery rejects concept pluralism precisely because he does *not* think the evidence bears out this prediction.

Let us now compare Machery's scientific eliminativist argument concerning concepts with my own concerning pain and pain "types."

Machery's formulations of both types of scientific eliminativist arguments imply that there are *some* successful generalizations that hold about the kind in question. For a vertical

argument, it is then argued that the generalizations are inconsistent because they are actually underwritten by *multiple* kinds with conflicting properties. For a horizontal argument, the successful generalizations are underwritten by either supersets or subsets of the kind, and no generalizations beyond those by which we identify the kind are *not* so underwritten. For pain and pain "types," by contrast, I argued that there just *are* no useful (i.e., reliable, valid, etc.) generalizations about pains or pain "types," because token pain experiences are explained by the activity of multiple, idiosyncratically-weighted mechanisms. This suggests that the scientific eliminativist argument I have developed here is a horizontal argument since the non-natural kind thesis relies on 'pain' and terms for pain "types" not picking out *any* natural kind.

There are, however, many natural kinds *implicated* in pain experiences; many of the mechanisms whose idiosyncratic weights explain each token pain are usefully generalized about. As discussed, for instance, Tylenol has reliable effects on inflammatory mechanisms, and stimulation of A-β fibres has reliable effects on processing in the dorsal horn. The utility of these generalizations, however, does not rest on their use of 'pain' or terms for pain "types": they become more *unreliable* when 'pain' or terms for pain "types" are employed instead. Tylenol does *not* have a reliable effect on inflammatory *pain*: Inflammatory pain often persists despite ingesting Tylenol—as common sufferers are all too aware. Stimulation of A-β fibres does *not* have reliable effects on pain—even pains antecedently diagnosed as centrally mediated. Tylenol and A-β fibre stimulation *are* appropriate treatments in some particular cases, but this appropriateness is not reliably generalized to other instances, even of the constructed pain "type." There are, then, useful generalizations about numerous mechanisms *implicated* in pain experiences, and those mechanisms are themselves natural kinds. As such, we may conclude that 'pain' and terms for pain "types" actually pick out *multiple* natural kinds; i.e., that my argument

for the non-natural kind thesis is a vertical argument.

I suspect that the uncomfortable fit between my argument for scientific eliminativism and either of the forms of that argument presented by Machery is that he denies concepts natural kindhood, but grants it to fundamental *types* of concepts. These three types of concepts are taken by Machery to realize the purported functional role of concepts. He then argues that the types of concepts realize this function in sufficiently different ways that classifying them together as one natural kind is inappropriate. In the case of pain, however, there are no such realizers. There is no single mechanism that explains *any* token pain experience, whereas Machery thinks there are three fundamentally different kinds of mechanisms that explain three fundamentally different types of concept deployment. There are no successful generalizations about pain "types" in the way that Machery thinks some generalizations about types of concepts are successful.

Rather than shoehorn my argument for the non-natural kind thesis into either the horizontal or vertical forms Machery offers, it is enough to point out that my reasoning for the non-natural kind thesis fully captures the spirit of both types of scientific eliminativist arguments, if not the exact letter of either. 'Pain' and referring expressions for pain "types" are useless for scientific inquiry because they do not adequately support generalizations. Horizontal and vertical scientific eliminativist arguments are simply two ways of showing that a term is not useful for scientific inquiry. We should only go pluralist, by contrast, if references to the kind in the relevant generalizations are useful.

3. Weiskopf

Having compared my scientific eliminativist argument with Machery's, I turn now to Weiskopf's (2009) critique to see if a parallel critique successfully counters my position. Weiskopf argues

that there is a substantial body of successful empirical generalizations about concepts, and that 'concept' is thereby a useful term for classifying bodies of knowledge stored in long-term memory used by default in cognitive processing. Weiskopf, like Machery, follows Boyd's natural kind criterion, and they agree that concepts have no underlying properties common to all. Nevertheless, Weiskopf identifies two major points of disagreement.[165] First, Weiskopf holds, contra Machery, that the empirical data supports seeing the conceptual systems *as a whole* as adaptive and having "coordinated" operations. Second, Weiskopf optimistically believes, while Machery pessimistically denies, that a psychological theory of concepts that unifies their operations is forthcoming—i.e., that a unifying theory is "*per se* in the cards." Of the similarities and differences between the two views, Weiskopf summarily concludes:

> He [Machery] and I agree on roughly what it takes for something to be a kind: kinds are sets of entities that share some large body of properties because of a common underlying mechanism. Machery holds that evidence similar to that reviewed here shows that concepts have no such common underlying properties, and hence do not qualify as a kind ("The plurality of concepts," 26).

While Weiskopf is spot-on in his identification of the similarities and differences between his and Machery's views, his diagnosis of Machery's mistake is incorrect.

Weiskopf claims that: "Machery's point, in essence, is that pluralism undermines concepts as a theoretical category of interest to psychologists."[166] He takes this to be in contrast to his own claim that: ""... the fact that a class of entities is partially heterogeneous at one level of analysis does not preclude their being homogenous at a higher level of analysis."[167] As we have seen, however, Machery is quite clear that pluralism of the kind advocated by Weiskopf is sometimes appropriate. Weiskopf provides no textual evidence, and I think there is none,

---

[165] Weiskopf, "The plurality of concepts," 27.
[166] Ibid., 29.
[167] Ibid., 27.

supporting his quoted claim that Machery rejects that some natural kinds are heterogeneous at one level of analysis and homogenous at another. In the case of concepts, Machery thinks not only that three fundamental types realize them, but that there are no generalizations underwritten by concepts proper that are *not* underwritten by these three subclasses; he thinks that the homogeneity of the class of concepts proper is insufficient to underwrite successful generalizations. Concepts proper do not have enough "scientifically relevant" properties in common.

Weiskopf's general positive argument for concept pluralism is as follows:

> 1) Psychological theories posit that concepts play a certain causal/explanatory role in cognition.
>
> Therefore, 2) Concepts should be identified with whatever entities or structures best fill that role.
>
> 3) Several distinct kinds of psychological structures satisfy the causal/explanatory role of concepts.
>
> Therefore, 4) Concepts are constituted by several distinct kinds of psychological structures.

Notice that Weiskopf is explicit that concepts are appropriately characterized functionally. This is in keeping with his belief that one task of empirical science is to discover the mechanisms that explain functionally characterized kinds:

> Empirical investigation sheds light on the inner nature of functional kinds by uncovering the structures and mechanisms that enable them to function as they do. So in the case of concepts, the task is to discover precisely which structures are responsible for the various functions and effects delimited by a psychological theory of concepts. ("The plurality of concepts," 6.)

Members of the class of concepts proper are held by Weiskopf to have "significant explanatory and functional commonalities," and one task of empirical investigation is to discover the mechanistic explanations of this (common) functionality. The relevant functional

characterization is that concepts represent categories in thought, have the capacity to be combined into larger structures, and play a further role in categorizations. Weiskopf argues that this functional homogeneity is not undermined by differences in the three realizers of these functions identified by Machery. On the contrary, these realizers are just the uncovered "structures and mechanisms" explanatory of concepts' homogenous functionality.

Weiskopf goes on to cast the debate between nihilists (i.e., eliminativists) and pluralists in explicitly functionalist terms. As above, empirical inquiry begins with an appropriate functional characterization of a kind. Functional characterization in hand, "… there are only three possibilities: (1) *One* kind of entity satisfies the role in question, (2) a *variety* of different kinds of entities satisfy the role, or (3) *no* entities satisfy the role."[168] If possibility 1 is actualized, there is presumed to be no support for eliminativism of any kind. When possibility 3 is realized, Weiskopf seems to think that "straightforward" eliminativism is supported; if nothing satisfies the functional role of concepts, then there are no such things as concepts. This is the first identified kind of nihilism, and though I think this is too hasty, I set it aside.[169] The second possibility is taken to be the interesting case: "The issue is what to say about the intermediate case in which we have a variety of realizers."[170] One may say that when there is a plurality of realizers the kind exists but is non-natural. This he considers a "weaker" kind of nihilism and includes Machery's scientific eliminativism about concepts as one such. As above, while I think Weiskopf is wrong to construe Machery's argument as based *solely* on concepts' plurality of realizers, he is right that whether we should go scientific eliminativist in the face of concepts'

---

[168] Ibid., 6.
[169] This embrace of what Weiskopf calls straightforward eliminativism is too hasty for the same reason that Machery's was. We saw that discovering that nothing satisfies a description (functional or otherwise) does not yet tell us whether we should eliminate the referring expression or whether we should revise the description.
[170] Ibid., 6.

(mechanistically distinct) plurality of realizers is their central point of dispute.

Facing multiple realizers, the central question is whether there are any generalizations supported by the functionally characterized kind *over and above* those supported by the realizers. Machery's claim is that for concepts there are none. Weiskopf maintains that there *are*, and that these further generalizations justify retaining the term 'concept.' These generalizations are held to be "insensitive to the particular kind of concept being used."[171] In particular, concepts proper are held to support generalizations concerning the logical form of concepts proper (32), conceptual combination (33), common modes of acquisition (35) and long-term memory processes (37). I pass over the details of these generalizations. What is important to note for my purposes is that if there *are* generalizations underwritten by concepts proper *in addition* to those supported by their discovered realizers, then concepts are a natural kind: this s the core of Weiskopf's and Machery's disagreement.

The moral of the Machery/Weiskopf story is that the key question for whether pluralism or scientific eliminativism is appropriate for a functionally-characterized kind realized by multiple mechanisms is: Are there useful generalizations about that kind that we would lose by eliminating the term referring to it from scientific inquiry? If the answer is no, then the kind should be split and the initial overarching kind term should be eliminated from scientific inquiry. This is Machery's position on concepts. If the answer is instead yes, then we should adopt a pluralism about the kind that takes the multiple, underlying realizers to be sub-kinds. This is Weiskopf's position.

Notice now that the central question for whether pluralism or scientific eliminativism about a kind is appropriate remains the same even if, as I have argued is the case in pain, there is

[171] Ibid., 31.

no single mechanism that explains any token. Pains are not multiply realized by different, single, underlying homeostatic mechanisms on different occasions. Rather, pains are multiply realized by an idiosyncratic activity of multiple mechanisms on *each* occasion. Despite this, whether we ought to go pluralist or scientific eliminativist turns on whether or not pluralism would allow retention of useful generalizations that scientific eliminativism would forego.

4. Scientific Eliminativism for Pain and Pain "Types"

Are there any useful generalizations we would lose by eliminating 'pain' and terms used to refer to pain "types" from scientific inquiry?

First consider 'pain.' If there are useful generalizations that we would lose by eliminating 'pain', then we should accept a pain pluralism in favor of scientific eliminativism despite the mechanical differences across token pains argued for in the previous chapters. I contend that there are not any.

Consider the following generalization schemas for pain:

P1: Pain is caused by _____.

P2: Pain is relieved by _____.

P3: People in pain _____.

My claim is that there are no reliable ways of filling in the blanks *useful for scientific inquiry.* There are, that is, no useful scientific generalizations about pain that would be lost by eliminating the term 'pain' from further scientific inquiry. This may initially strike the reader, as it struck me, as highly unintuitive.

It initially seems as if P1-P3 can be usefully filled in quite easily. We may try, for instance, the claim that pain is caused by getting hit in the head. Notice, however, that pain is *not*

reliably caused by getting hit in the head—not reliably enough for any scientific generalizations. I will not feel pain by a *slight* hit in the head, and what counts as a *slight* hit in the head is idiosyncratic; it will vary from person to person and from circumstance to circumstance. The obstacles here are everyday variation that we routinely take for granted, but that nonetheless block useful generalizations *for scientific inquiry*.

The central goal of inquiry in pain science is *treatment*; treatment is what we want our generalizations about pain to be reliable *for*. If we try to fill in the blanks for P1-P3 such that the generalizations are reliable for *treatment* purposes, we fail miserably. We may, for instance, try filling in P2's blank, as initially seems plausible, with 'morphine.' Thus, we may claim, pain is relieved by morphine. This generalization, however, is not enough for *treatment* purposes. The *amount* of morphine is crucial; too little and the pain is unrelieved, too much and the patient dies—setting aside that for some "types" of pain, morphine is an altogether inappropriate treatment. Notice that how much morphine *you* need, a crucial quotient for treatment purposes, is idiosyncratic. Treatment is idiosyncratic not only to the individual, but to the same individual at different times. If we assess the utility of generalizations about pain proper for treatment purposes, none will be successful.

Is the situation different for pain "types"? Consider the following generalization schemas:

PT1: Inflammatory pain is caused by _____.

PT2: Inflammatory pain is relieved by _____.

PT3: _____ is a reliable treatment for chronic pain.

PT4: _____ is a reliable treatment for headaches.

I admit that we can come *closer* to filling in these blanks reliably for pain "types" than for pain proper. PT2 is *often* true if 'Tylenol' is put in the blank. Generalizations for some pain "types"

are *more* reliable than for pain proper because many pain "types" *are more* natural than pain proper.

Recall from chapter 1 that the naturalness of a kind is a matter of degree that depends on the success of the generalizations about that kind. Pain "types," recall from the previous chapter, are most often constructed by etiology or symptoms. Fewer mechanisms are implicated in these pain experiences than the greater number of mechanisms implicated by the class of *all* pain experiences. It is therefore expected that these pain "types" are *more* natural than pain proper; it is to be expected that degrees of naturalness will differ relative to how many mechanisms a given "type" of pain implicates. Inflammatory pain comes much *closer* to identifying the activity of a single mechanism than neuropathic pain, with headaches (most likely) being somewhere in between. It is to be expected, then, that generalizations about inflammatory pain are *more* reliable, since they implicate fewer mechanisms than generalizations about neuropathic pain. Our ability to fill in PT2 with something that is *often* true contrasts with our utter loss at what to fill in for PT3. Fewer (though there are always many) implicated mechanisms for a pain "type" increases the reliability of generalizations about it because there is less possible variation in the idiosyncratic convergence of the mechanistic activities.

These comparative considerations may lead us to conclude that pain "types" *are* natural kinds that we just do not know enough about yet: we know more about the mechanisms for some types of pain than others, and eventually all pain "types" will be mechanistically explicable. This conclusion, however, would be in error.

Closer consideration reveals that even inflammatory pain, the most natural of the pain "types" schematized above, is nevertheless not a natural kind. No useful generalizations would be lost by eliminating the term 'inflammatory pain,' so long as we retained terms for the

inflammatory *mechanisms* implicated in "inflammatory pain" instead. Any success of generalizations about "inflammatory pain" is a function of the naturalness of these implicated inflammatory mechanisms. Notice that not only would we not *lose* any utility, but some would be *gained* if we replaced any tempting generalizations developed from the schema in PT1 or PT2 with generalizations about inflammatory *mechanisms*. If we generalize only about the implicated mechanisms, then we can reliably characterize causal relations—including the interventions so important for treatment.

As soon as we include the word 'pain,' we will *lose* some utility due to the idiosyncratic activity of the multiple mechanisms underlying each pain experience. Inflammatory pain is often helped by Tylenol—but not always. So too, as with pain proper, how *much* Tylenol is required is idiosyncratic. How much more reliable, how much more useful, are generalizations about the mechanism implicated in inflammatory pain. Even as these things are true of inflammatory pain, they are more stark for even *less* natural kinds like headaches and chronic pain. So *many* mechanisms are implicated in these "types" of pain that the causes and treatments are even more idiosyncratic.

It is important to remember that it is the kinds that matter, and not the terms we use to refer to them. One could simply *call* all C-fibers firing 'C-fibre pains.' As we can make reliable, useful generalizations about C-fibers, it would then follow that "C-fibre pains" were natural kinds. None of this, however, would illuminate *pain*. As discussed extensively in the previous two chapters, C-fibre activation is not co-extensive with anything we would antecedently characterize as pain: it is not reliably correlated with the sensation of pain, the perception of tissue damage, or motivations to withdraw or avoid, *much less* all of these aspects together as is present in a paradigmatic pain experience. We can *call* all C-fibres firing a type of pain, and

conclude that pain "types" (at least this one) are natural kinds. The claim is obviously uninteresting. Sets can be constructed and labeled any way one likes. My claim is that there are no types of *pain*, as we antecedently characterize pain, about which there are reliable, scientific generalizations.

5. Conclusion

The moral of the Weiskopf/Machery debate is that whether pluralism or scientific eliminativism about a kind is appropriate is determined by whether there are useful scientific generalizations about that kind. For concepts, Machery and Weiskopf agree that conceptual deployment is realized, on different occasions, by different mechanisms, but disagree about whether those differences are *enough* to undermine utile generalizations about the class of concepts proper. They disagree, in short, about whether utility would be lost by eliminating reference to the class of concepts proper and retaining references only to its multiple realizers for scientific inquiry.

In the case of pain and pain "types," I have argued that there is *not* enough commonality to underwrite utile scientific generalizations about them and that no utility would be lost by retaining terms referring only to the many mechanism they implicate. The differences in the mechanistic explanations of different token pain experiences blocks useful generalizations about pain or pain "types"; the differences are non-negligible for scientific inquiry. Scientific eliminativism about pain loses no generalizations useful for *treatment,* and it is treatment above all that is the goal of pain science. Insofar as scientific generalizations about pain or pain "types" *are* successful, the generalization becomes even *more* reliable when the pain-related terms are replaced by terms referring to the mechanisms implicated by the pains.

Though it is admittedly more intuitive than the non-natural kind thesis, pluralism about

pain should be resisted in favor of a scientific eliminativism. In the concluding chapter, I argue that despite this, we should *not* be what Machery calls traditional eliminativists about pain. Though eliminating the term 'pain' and terms for "types" of pain in scientific inquiry would not result in a loss of utility, eliminating these terms from daily life *would*. Though scientific eliminativism for pain ought to be embraced, traditional eliminativism, like pluralism, is appropriately resisted.

Chapter 6: Eliminativism

1. Introduction

In the previous chapter I argued that pluralism about pain should be resisted because the differences in the mechanistic explanations of token pain experiences undermine the utility of reference to pain or pain "types" in the generalizations of any relevant science. Though many of the mechanisms implicated in pain experiences *are* usefully generalized about, none would be lost by eliminating the term 'pain' or terms for pain "types" from scientific inquiry. Similar reasoning may be employed to further argue that pains do not exist, and that reference to pain or pain "types is as useless for everyday life as I have argued it is for scientific inquiry. That a term never successfully refers and should to be eliminated from everyday speech is a position I will call *folk* eliminativism. In this chapter, I argue that folk eliminativism for pain is as inappropriate as pluralism.

The folk eliminativist arguments currently on offer in the literature are a helpful starting place for exposing the relevant issues. I turn first to Dennett's eliminativism in section 2, and argue that his offered arguments support the claim that pain is not a natural kind, but do not establish folk eliminativism. Once the folk notion is relieved of the inappropriate burdens Dennett places on it, its successful use in everyday discourse is no longer undermined. I turn in section 3 to Hardcastle's eliminativism. I argue that not only is Hardcastle's folk eliminativism in tension with her positive proposal of the nature of pain, but that both positions are undermotivated.

Though both of these influential arguments for folk eliminativism fail, I argue in section 4 that my own arguments for scientific eliminativism about pain do not support folk eliminativism. I have developed the view that each token pain is explained by an idiosyncratic

convergence of the activity of multiple mechanisms. While this idiosyncrasy blocks useful scientific generalizations about pain, referring to those idiosyncratic experiences is useful for everyday life. Pain is very real and very important. That neither pain nor any "type" of pain is a natural kind does nothing to undermine this truth. I conclude with some general reflections on how the dissertation as a whole can be understood as a case study in how a kind that is usefully referred to in daily life may nonetheless be uselessly referred to in scientific generalizations.

2. Dennett's Eliminativism

Dennett's eliminativism about pain is characteristically subtle, and he offers two general, and related, arguments for it. The first relies on his distinction between personal and sub-personal levels of explanation. He argues that pain is a personal, non-mechanistic, phenomenon, and as such is not appropriately identified with any of the sub-personal, mechanical, events or processes in the brain. Since it is to the mechanical, scientific categories that Dennett maintains we are to look for our best ontology, he concludes that 'pain' never successfully refers. His second argument focuses on the specific reason the non-mechanical, person-level phenomenon of pain cannot be identified with any mechanical, sub-personal processes—namely, and we cannot have a true theory of pain even at the *personal* level. Our ordinary concept of pain is argued to be riddled with deeply entrenched contradictions. Since no true theory is contradictory, Dennett concludes that there is no true theory of pain. I examine the first argument before turning to the second.

In *Content and Consciousness* and in "Personal and Sub-personal Levels of Explanation," pain serves as Dennett's example of a personal-level phenomenon. Persons feel pain; pain is not experienced by a *part* of the person, nor is it rightly identified with any set of functions occurring

within the person. Together, these facts are presumed to block inquiry into the *mechanics* of pain, and it is mechanical questions that are raised, and answered, at the sub-personal level:

> An analysis of our ordinary way of speaking about pains shows that no events or processes could be discovered in the brain that would exhibit the characteristics of the putative 'mental phenomena' of pain, because talk of pains is essentially non-mechanical, and the events and processes of the brain are essentially mechanical (*Content and Consciousness*, 91).

Dennett takes the folk to be committed to pain being a brute and unanalysable experience. Both *how* we discriminate pains as painful and *what* we discriminate as painful are simply brute facts that do not submit of a mechanistic explanation.[172] These features of pain-talk are so well entrenched, according to Dennett, that if we abandon them, we cease to talk about pain at all. Consequently, we can neither deny that pain is unanalysable, nor provide the mechanical theory of pain that requires this:

> Since the introduction of unanlysable mental qualities leads to a premature end to explanation, we may decide that such introduction is wrong, and look for alternative modes of explanation. If we do this we must abandon the explanatory level of people and their sensations and activities and turn to the *sub-personal* level of brains and events in the nervous system. But when we abandon the personal level in a very real sense we abandon the subject matter of pains as well. When we abandon mental process talk for physical process talk we cannot say that the mental process analysis of *pain* is wrong, for our alternative analysis cannot be an analysis of pain at all, but rather of something else—the motions of human bodies or the organization of the nervous system. … Abandoning the personal level of explanation is just that: *abandoning* the pains and not bringing them along to identify with some physical event (*Content and Consciousness,* 94).

We cannot provide a description, ostention, or mechanical explanation of how we discriminate pain it without *abandoning* the ordinary concept of it. The ordinary concept of pain, therefore, does not admit a mechanical explanation.

Do pains then not exist? Dennett is characteristically cagey about the ontological question.

---

[172] Dennett, *Content and Consciousness*, 92-93.

He is clear that "Pains are feelings, felt by people, and they hurt. People can discriminate their pains… we do talk about the qualities of sensations and we act, react and make decisions in virtue of these qualities we find in our sensations."[173] This statement concerning what pains are, however, is Dennett's conceptual analysis of the *folk* concept of pain, and folk, or ordinary, concepts are not taken by Dennett to be our best guide to ontology. Scientific theory is. If we take the goings on in the brain as providing our best ontology for mental terms as we ought, then folk mental terms used to refer to a phenomenon we *cannot* identify with goings on in the brain do not exist. Thus he says:

> …the *terms* [intentional, mental terms] are construed to be in different categories and only one category of terms is referential. The only way to foster the proper separation between the two levels of explanation, to prevent the contamination of the physical story with unanlysable qualities or 'emergent phenomena', is to put the fusion barrier between them. Given this interpretation it is in one sense true that there is no relation between pains and neural impulses, because there are no pains; 'pain' does not refer. There is no way around this. If there is to be *any* relation between pains and neural impulses, they will have to be related by either identity or non-identity, and if we want to rule out both these relations we shall have to decide that one of the terms is non-referential ("Personal and Sub-Personal Levels of Explanation," 21).

The "fusion barrier" alluded to here allows for the sort of humuncular functionalism Dennett advocates elsewhere.

Dennett's humuncular functionalism prescribes that we begin theorizing about the mind with the intentional, mental terms of the folk, and break down these mental items into increasingly compartmentalized functional components as guided by scientific inquiry. Each step results in evermore compartments that are more "stupid" than the last—that is, increasingly less intentional or mental. Employing intentional idioms, however, is a crucial part of investigation and useful for the development of scientific models (e.g., psychological or neuroscientific).

---

[173]Ibid., 94.

Despite this (at least initial) utility, intentional idioms give way to some mechanistic, non-intentional explanations:

> There is no reason why intentional terms cannot be used provisionally in the effort to map out the functions of the behavior control system of men and animals, just so long as a way is found eventually to 'cash them out' by designing a mechanism to function as specified... (*Brainstorms,* 62).

That the models of mentality must ultimately be mechanical seems to be taken by Dennett as simply obvious:

> This much ought to be uncontroversial; though psychologists may ignore the details of realizations while elaborating and even testing their models, the model-making is ultimately bound by the restriction that any function proposed in a model must be physiologically or mechanically realizable one way or another (*Brainstorms,* 41).

That the intentional idioms, first taken from the folk, are allowed as an appropriate part of the *development* of these mechanical models is a testament to the *fusion* between the personal and sub-personal levels of explanation. Mental/intentional terms are "fused*"* into non-intentional, mechanistic terms in the light of scientific inquiry.

Dennett's eliminativism about pain, then, is not based *solely* on the fact that it is a personal-level phenomenon. Some personal-level phenomenon and the terms that refer to them are usefully "fused" into the mechanical, functional explanations provided by scientific inquiry. The problem with pain is that the personal-level phenomenon will *never* submit to a mechanical explanation without our changing the subject. And again: since we should take our ontological cues from those mental items that *will* submit to a mechanical explanation, those that will not, do not exist. Hence, pains do not exist.

Dennett's first argument for pain eliminativism is then reconstructed as:

> 1) The ordinary concept of pain is, essentially, as of an unanalyzable mental quality, brutely discriminated.

2) Unanalyzable mental qualities that are brutely discriminated will never submit of a mechanical explanation.

3) Therefore, a mechanical explanation of pain is not possible.

4) Real entities admit of mechanical explanation.

5) Therefore, pain is not a real entity.

I have argued for conclusion 3 in the previous chapters on other grounds, and would put this important claim by saying that pains are not natural kinds. Insofar as the non-mechanical nature of pain is taken to support the conclusion that pains therefore *do not exist*, however, Dennett errs. Premise 4 is false. I argue against inferring non-existence from non-naturalness in section 4.

For now, however, notice that even if Dennett's move from the non-naturalness to the non-existence is accepted, premise 1 is objectionable. Why should we construe the folk as being committed to the claim that pains are essentially unanalyzable mental qualities that are brutely discriminated? Dennett does not offer any reasons. We should, at the very least, be very suspicious of these claims, since 'essentiality,' 'mental quality,' 'unanalyzable,' and 'brutely discriminated,' all reek of the vocabulary of the specialized theorist and not the layman. Dennett's argument turns on this claim about the folk concept that receives no support.

By Dennett's lights, however, the problem with the folk notion of pain goes beyond its not identifying a mechanically explicable phenomenon; the folk notion of pain is contradictory. Though also resting on a contentious interpretation of the folk, this argument is more important than the first.

Dennett presents his argument that the folk notion of pain is irresolvably contradictory, and that therefore its referent does not exist in "Why You Can't Make a Computer that Feels Pain." I think Dennett is right that *if* an ordinary concept *were* contradictory beyond revision it

should be eliminated. I deny, however, that the folk concept of pain *is* contradictory in the ways Dennett alleges, much less unrevisably so. His arguments along the way usefully support that pain is not a natural kind, but they do not support the further claim that pain is not a usefully referenced *folk* kind.

Dennett's argument for the contradictory nature of the folk pain-concept is developed as a failed, attempted response to the "skeptical" claim that it is impossible to make a computer that feels pain. Having explained at length why behaviors, both internal and external, will never satisfy the skeptic, Dennett's strategy for refuting the skeptic is: "… to set down the bare bones of current physiological theory relating to pain, and list some of the attested pain phenomena, the data any acceptable theory of pain must account for, and see if there are any insuperable difficulties [for a computer instantiating the theory] presented by them."[174]

By supplementing the gate-control mechanism, he develops a flow chart meant to serve as our best sub-personal explanation of pain—a mechanical model of our best theory.[175] He argues that affixing the label 'pain' to any given box or adding an additional one anywhere in the chart would be contentious and, more importantly, add nothing explanatory. Of course, the flow-chart was never intended to include pain as a box *within* it; rather, the whole complex system of interacting mechanisms was supposed to explain pain.

The problem, however, is that the multiple mechanisms presented in the flow chart dissociate:

> The chief value of all this somewhat science-fictional flow-charting and compiling of odd phenomena—the reason I have spent so much time on it—is that it serves to drive wedges of contingency between features that are often

[174] Dennett, *Brainstorms,* 199.
[175] Throughout the article Dennett seems to think that Melzack and Wall's gate-control theory was intended to provide an exhaustive account of *the* pain mechanism. As I have argued in chapter 3, Melzack and Wall's model was never intended for that purpose.

> thought to be conceptually inseparable, simply because they are usually coincident. What I am asserting is that the arrows on the flow-chart are the arrows of normal causal relation, and wherever we have seen fit to posit a particular relation or dependency, we can imagine a severing of the normal connections responsible for it. Some of this fragmentation has familiar manifestations, some is to be found only rarely, and some never occurs, so far as I know, though we can conceive of it occurring. (*Brainstorms*, 220.)

Pain is a complex experience, each dimension of which may occur in the absence of the others. This is the same problem I raised for unitary accounts of pain in chapter 2. An adequate theory of pain, according to Dennett, should determine which combination of activations result in pain and which do not.

If accurate, the model Dennett provides would identify all the mechanisms implicated in pain experiences, but whenever some of these mechanisms are active independently of the others, the model cannot be used to determine whether the resultant experience is pain. It is useless for the legion dissociative cases. According to Dennett, however, this is what we would want a *philosophical* theory of pain to do; an adequate philosophical account of pain should allow us to say of each token experience whether it is pain or not.

For principled rulings on the dissociative cases, Dennett thinks the philosopher must determine what is *essential* to pain. To determine what is essential to pain, furthermore, Dennett holds that we must consult (well-entrenched) folk intuitions as evidenced by common usage of 'pain.' Enter the alleged problem: Ordinary, folk intuitions about pain comprise an inconsistent set.

Dennett presents three contradictions in the ordinary, folk concept of pain. He suggests unproblematic ways of dealing with the first and somewhat problematic ways of dealing with the second. It is the third contradiction, however, that he takes to support his eliminativism because it is irresolvable. I consider each in turn.

Here is the first:

- Contradiction 1: According to ordinary usage: 'Pain' is properly used as both a sortal *and* a mass noun.

'Pain' does seem to be sometimes used as a sortal—there is a pain in my foot and another pain in my arm—and as a mass noun—I am in *more pain* today than I was yesterday. The superficial difference in the grammar of these two uses of 'pain' may indeed merely be superficial, or we develop two conceptions of pain. After all, there are other terms embodying the same "contradiction" that we do not find very problematic. Consider 'chicken': "There is *a* chicken on my plate," and "There is *more* chicken on your plate than mine," are both perfectly acceptable uses of 'chicken.' This superficial grammatical difference does not undermine our confidence in the existence of chickens. Dennett is open to an unproblematic revision of this first contradiction if necessary.

Here is the second:

- Contradiction 2: According to ordinary usage: For pain *esse is percipi* and pains can occur unperceived.

This second contradiction is taken to be more entrenched than the first, but revision is nonetheless available. Dennett thinks that theoretical considerations favor revising in favor of the intuition that pains can occur unperceived as against the intuition that for pain *essi is percipi*. Dennett seems to think that we would do better to build our theory of pain based on the idea that pains are the result of a perceptual process since, he says, they obviously normally are—and all perceptual processes sometimes occur unperceived.

Taking this side of the clash resolves the second contradiction, but Dennett thinks it leads us to the irresolvable third, since:

> [p]erception has not one product but many, operating at different levels and in different ways…. Looking at all the various effects such an interpretation of

> signals could produce, we can answer the philosophical questions about pain only be deciding which effects are 'essential' to pain and which are not (*Brainstorms,* 226).

Dennett claims that we must call upon yet further, well-entrenched intuitions to decide which products of the perceptual process are essential to pain and which are not. It is *these* intuitions that cannot be revised into harmony. Thus:

- Contradiction 3:

A) Pains are essentially items of immediate experience or consciousness; the subject's access to pain is privileged or infallible or incorrigible.
B) Pains are essentially abhorrent or awful.

Dennett maintains that "[e]fforts to capture both of these 'essential' features in a theory of pain are bound to fail…."[176] To make the contradiction transparent, he paraphrases each half so that they wear their modalities on their sleeves. Thus A becomes:

> A*) It is a necessary condition of pain that we are 'incorrigible' about pain; i.e., belief that one is in pain is a *sufficient condition* of pain.

While B becomes:

> B*) It is a necessary condition of pain that pains are awful.

From A* and B* it follows that believing one is pain is sufficient for experiencing something awful.

The problem is that the sufficiency of A* and the necessity of B* result in contradictory rulings in the cases of non-awful pains. Note again that nothing in the mechanical model of pain can resolve the contradiction. Dennett rightly points to dissociative cases, like those in chapter 2, of one believing that they are in pain and that they are *not* experiencing something awful. Not only do those undergoing such experiences report them this way (e.g., morphine, lobotomy, pain asymbolia), but the folk accept these reports as veridical. Dennett claims that some serious

---

[176] Ibid., 226.

conceptual back flips would be required to harmonize the folk intuitions in these cases; e.g., perhaps one does not believe what one believes that one believes.

Rather than attempt any such "heroics," Dennett advocates giving up not only these contradictory intuitions, but all of the folk intuitions about what is essential to pain:

> I recommend giving up incorrigibility with regard to pain altogether, in fact giving up *all* 'essential' features of pain, and letting pain states be whatever 'natural kind' states the brain scientists find (if they ever do find any) that normally produce all the normal effects (*Brainstorms,* 228).

Read in isolation, this passage suggests that Dennett is open to the possibility that pain (brain) scientists will eventually arrive at a true theory of pain, but this is not so. According to Dennett, any true theory will not be the true theory of *pain*—though it may be a theory of something else that the folk will later begin to *call* pain. When we give up the (contradictory) essentiality claims about pain in favor of whatever the scientists discover, we give up the folk notion of pain. According to Dennet, to give up those intuitions is to forsake belief in the existence of a referent of the (current) folk concept of pain:

> …if, as I have claimed, the intuitions we would have to honor were we to honor them all do not form a consistent set, there can be no true theory of pain, and so no computer or robot could instantiate the true theory of pain, which it would have to do to feel real pain (*Brainstorms,* 228).

The skeptic triumphs: Computers will never instantiate a true theory of pain because there is not one.

Dennett resists the conclusion that the third alleged contradiction also shows that *humans* do not feel pain:

> The parochiality of the concept of pain protects us but not robots (or Martians or at least lower animals) from the skeptical arguments, by fixing the burden of proof: an adequate theory of pain must have normal human beings as instantiations, a demand that presupposes the primacy, but not the integrity, of our ordinary concept of pain (*Brainstorms,* 228).

That humans feel pain is an essential part of the very concept of pain, even if that concept is also contradictory. It is *not* an essential part of that concept that computers experience pain. To evaluate whether *computers* ever (or ever could) feel pain we must evaluate a proposed theory. The contradictory nature of the folk concept of pain will guarantee that the skeptic can always point to an intuition that the theory instantiated by the computer does not capture. When evaluating whether *humans* ever feel pain, however, we need only consult the well-entrenched intuitions of the folk to learn that they do.

Dennett's eliminativism about pain, then, rests on his belief that a true theory of pain encompass the well-entrenched folk intuitions, but their contradictoriness makes this impossible. On the contrary, however, the folk conception of pain is *not* contradictory on the simple grounds that the folk do not have the well-entrenched intuitions that Dennett claims.

Consider the second contradiction: The essence of pain is perception, and pains can occur unperceived. Why believe that the folk are committed to this essentiality claim? Indeed, the contradiction is supposed to arise for Dennett because there are some cases that the folk are willing to countenance as pain that are *not* the result of any perceptual process. Dennett wants to revise to exclude any pains that are *not* the result of a perceptual process since pains normally *are* the result of a perceptual process. Why, however, should we go beyond this "normally*"* and attribute to the folk the claim that pains are *essentially* the result of a perceptual process. The folk concept of pain, I have argued, is of a complex experience that paradigmatically includes sensation, perception, emotion, cognition *and* behavioral responses. Why construe the folk as holding that *any one* of these components is either necessary or sufficient—including the perceptual component? Why attribute any essentiality claims to them? We should, instead, simply reconcile the allegedly contradictory intuitions by interpreting the folk concept of pain as

of an experience that *paradigmatically includes* perception, but does not always. If we do, the supposed contradiction, and therefore the folk eliminativism on these grounds, is avoided.

Similar considerations apply to Dennett's third supposed contradiction. Why attribute to the folk the claim that pains are *necessarily* awful when, as
Dennett discusses himself, there are dissociative cases of non-awful pains that the folk still report and countenance as pain? Here again, if the folk concept of pain is of a complex experience with multiple components no one of which is necessary or sufficient, then contradiction is avoided. Pains are *paradigmatically* awful, but they need not be. If the folk are not committed to the claim that pains are *essentially* awful, as I have argued in chapter 2 that they are not, then this third contradiction provides no support for an eliminativist conclusion.

Dennett goes wrong in thinking the folk concept of pain is inherently contradictory, but this should not eclipse what Dennett has gotten *right*. The folk concept of pain will *not* submit to a mechanistic explanation and for reasons similar to those offered by Dennett. The folk concept of pain implicates multiple, dissociating mechanisms. This problem of dissociation is exacerbated, as I argued in chapter 4, by the idiosyncrasy of the interactions of the dissociative mechanisms that explain any token pain experience. As such, 'pain' is not a useful term for doing science. Furthermore, Dennett is right that philosophers have attempted to clean up the notion of pain to provide its necessary and sufficient conditions to the detriment of their theories. This is to be expected if, as I've argued, the ordinary, folk concept of pain is of a complex paradigm that tolerates a wide range of variation. Dennett himself comes close to appreciating this leniency in the folk concept:

> This … slackness in our ordinary use has provided a playground for interminable philosophic disputation over the issue: can there be unfelt pains? I suggest that our flow-chart handles this traditional question by discrediting it. There can be, in principle, any combination of the normal 'causes and effects' of pain in the

> absence of any others, and intuitions will no doubt clash about which words to use to describe the results (*Brainstorms,* 222).

The ordinary concept of pain does have a lot of "slack," and this looseness in the folk notion *has* given rise to endless philosophical controversies about how the loose notion should be firmed up. Intuitions *will* vary in dissociative cases and they will vary all the more the farther from "normal" (i.e., paradigmatic) that we get. Most important, Dennett is right that scientific theories of pain show that pain does not submit of a mechanistic explanation. From this we should conclude that pain is not a natural kind. Insofar as Dennett argues that pain does not submit of a mechanistic explanation and that the term therefore gets no purchase for *scientific* theorizing, I wholeheartedly agree.

Dennett goes wrong by committing the same mistake he aptly diagnoses in others; he tries to remove the slack from the folk concept by illegitimately saddling it with essentiality claims. That slack is here to stay. There is no good reason, and none offered by Dennett, to think that the folk concept of the pain is committed to any one component of the complex experience of pain being either necessary or sufficient. Dennett's arguments that the folk notion of pain is contradictory and that therefore we should be eliminativist about pain therefore fails. Whether an argument for folk eliminativism based purely on *scientific* eliminativism can succeed even after these essentiality claims are rejected will be considered in section 4.

3. Hardcastle

In *The Myth of Pain*, Valerie Hardcastle presents arguments for both a positive view of pain's nature and elimination of the folk notion of pain, and in this section I present arguments against both. The positive view is a dual-systems model according to which humans have a pain system comprised of two independent and distinct systems: the pain sensory system and the pain

inhibitory system. Her eliminativism prescribes that we eliminate the folk notion of pain—characterized as a simple sensation that is necessarily conscious and necessarily awful. It will be helpful for establishing why her folk eliminativism is unsuccessful to spend a considerable amount of time on her positive account.

I agree with Hardcastle on two central points. First, Hardcastle's wide canvas of the empirical literature on pain leads her to conclude that (pain exists and) is a complex, multi-dimensional phenomenon. Second and related, she accurately sees that most philosophical accounts of pain have failed because they have tried to identify pain with some one dimension of its complexity. Aside from the implication that there is a pain-specific "processor," Hardcastle herself perhaps best summarizes these points of agreement:

> Neither our conscious experience of pain, the damaged tissue itself, nor our bodily or emotional reactions are fundamental to pain processing. Each is but one component of a larger processor. Hence, it is a mistake to try to claim one or the other as pain simpliciter. And it is equally erroneous to conclude that since we cannot identify one or the other with pain, there is no such thing (*The Myth of Pain*, 146).

Unfortunately, however, these two key points are where agreement ends.

Ultimately, Hardcastle characterizes pain as being akin to any other sensory modality. Her theorizing about the nature of pain, accordingly, embraces a traditional dual model of pain processing that is kicked off by the activity of nociceptors. These primary afferents transmit chemical, mechanical and thermal information via two branching streams, both of which terminate in cortical processing, and all of which is subject to feedback loops. This system she calls the pain sensory system, or PSS:

> The entire pain sensory system functions largely in the same way as any of our sensory systems. Their pieces are united by our best guess of their function… The components of our auditory system take the information contained in atmospheric compression waves and use it to compute the placement of things. And the components of our pain system take pressure, temperature, and chemical readings

> of our surface (and interior) and use this information to track what is happening to our tissues. The A-δ cells and the C-fibers do this, as do the spinothalamic tract and its connections to cortex. In sum, it appears that we have a complex but well-defined sensory system that monitors our tissues to promote the welfare of our bodies. (*The Myth of Pain*, 107.)

Before turning to Hardcastle's need to supplement the PSS to explain pain experiences, two preliminary problems with this model should be noted. First, the identified stimuli here are pressure, temperature and chemical information. As discussed in chapter 2 and supported by the pain research discussed in chapters 3 and 4, this will not characterize a distinct modality—or "sensory system"—for pain as against touch. Second, the model of the PSS is illegitimately over-simplified. In additional to placing undue weight on the traditional "nociceptors," it identifies far too few relevant pathways and synapses.

It is a bit odd that Hardcastle presents the model as confidently as she does, since she acknowledges that it does not do justice to the complexities of pain–relevant processing. Despite acknowledging that, "[a]t the least, it is clear that the classic view of pain processing is woefully oversimplified and quite limited,"[177] she retains her further *weakened* version of the traditional model anyway. Quotes like this one and her acknowledgement that the scientific literature posits elements of nociceptive processing (much less *pain* processing) unaccounted for by her PSS make it unclear whether we should take it too seriously—or whether she takes it very seriously herself.

More important than its inadequately capturing nociceptive processing, Hardcastle recognizes that the PSS model does not adequately capture *pain.* The PSS is, at best, a model of nociceptive processing and pain is not well correlated with nociception. In this way, pain is admitted to be *unlike* our other sensory systems:

---

[177] Ibid., 108.

> Our other perceptual systems do not work in this fashion [yielding perceptual experiences that are poorly correlated with the objective phenomena they are supposed to be giving us information about]. There is a highly reliable correlation between having the rods and cones in our retina being bombarded by light photons and having some visual experience or other. … Any theory of pain is going to have to explain why our peripheral sensors for noxious stimuli are not well connected to our sensations of pain. Indeed, the relation between external events and internal indicators is part of what individuates our systems. Without better correlation between external events and internal activity, we cannot claim that our putative pain system is in fact a pain system. (*The Myth of Pain*, 124.)

Since pain experiences are not well correlated with activity in the processing system sketched by Hardcastle, why consider that processing system the *pain* processing system at all?

Hardcastle's answer: The PIS. The PIS is posited to explain not only the lack of correlation between nociception (activity in the PSS, she thinks) and pain experiences, but the subjectivity one may be tempted to think true of pain *because* of its poor correlation with nociception:

> … we actually have two separate systems involved in our perceptions of pain. One functions as a pain sensory system (PSS), quite analogously to our other sensory systems. The other (PIS), which developed independently of our PSS, actively inhibits its functioning.
> Differentiating between the two systems helps explain the remaining controversies surrounding the basic nature of pain. While a PSS supports a perceptual view of pain as a completely objective phenomena, adding in a PIS (without explicitly recognizing this fact) accounts for the strongly subjective aspects of pain. I shall claim that a PSS functions according to the same basic rules of all our sensory systems and insofar as the pain system is a simpler system than, say, vision or audition, it makes sense to take pain perception as a paradigm instance of a conscious experience. But, insofar as we also have a PIS, pain also becomes a special case in our collection of conscious phenomena. (*The Myth of Pain*, 96.)

Because of the PSS, pain is like every other sensory system, but because of the PIS it is a sensory system unlike any other. Hardcastle regularly emphasizes these two sides of pain's coin; e.g., "I argue that our pain sensory system functions according to the same basic rules of all our sensory

systems,"[178] but later, "There are several important empirical facts that any theory of pain needs to be able to explain. These are facts that, by and large, do not have analogues in our other perceptual systems."[179] The PIS is intended to let Hardcastle maintain a perceptual view, while still acknowledging that pain experiences are not well correlated with activity in the identified sensory processing system (PSS). Whenever the pain system is unlike other sensory systems (e.g., the complexity, the involvement of almost the entire brain, the seeming subjectivity, and the poor correlation between pain experiences and nociception), the posited PIS is supposed to explain the difference.

Hardcastle's first pass at establishing the existence of the PIS is to say that it is a system that is anatomically distinct from the PSS. To support this claim she writes:

> We have known for some time that many of the inhibitory streams are not merely feedback loops in our ascending nociceptive fibres, for they are anatomically distinct from these processors. Three areas are primarily responsible for inhibiting nociceptive information in the spinal column: the cortex, the thalamus, and the brain stem. The dorsal raphe is probably heavily involved as well. In particular, neocortex and hypothalamus project to the periaqueductal gray region (PAG), which then sends projections to the reticular formation… The reticular nuclei then work to inhibit activity in the dorsal horn. This processing stream works by preventing a central cortical representation of pain from forming. … this inhibition stream does not merely disrupt the transmission of nociceptive information, but it actively prevents it from occurring. (*The Myth of Pain*, 132.)

This is a very curious passage in Hardcastle for two main reasons.

First, many of the anatomical areas she identifies as proprietary to the PIS are identified as relevant areas in her own proposed model of the PSS. In what sense are the two posited systems then anatomically distinct? Perhaps there are other neuronal pathways flowing through these same areas, but that evidence needs to be (and is not) offered. Even then we would not have distinct anatomical areas—only distinct neural circuits throwing through the *same*

[178] Ibid., 94.
[179] Ibid., 122.

anatomical areas. Moreover, areas crucial for the proposed PIS *not* included in her own proposed model of the PSS, particularly the PAG and the hypothalamus, *are* included in other dominant models of nociceptive processing.[180]

Second, and equally troubling, is the emphasis placed on the important role of downstream cortical processing in the PIS. This is problematic for *Hardcastle* because her second of two attacks on gate-control theories is her denial that appealing to such processes is legitimate.

Hardcastle attacks gate-control theories of the Melzack-Wall variety for two main reasons that both apply directly to her own dual-systems view.[181] First, she says, we "…do not have anything approaching a complete theory at this stage in the game."[182] The incompleteness of the gate-control theory is offered as a decisive criticism:

> Though the [gate-control] theory does explain low-level phenomena, nothing has been worked out in particular regarding the clearly "psychological" influences. Melzack and Wall themselves simply gesture toward a central gating mechanism that presumably would explain hypnosis effects, any remaining chronic pains, the dismal correlations between ER injuries or back injuries and pain, SIB [self-injurious behavior], phantom limb pain, and so on. But without providing more details—where this central gating mechanism is located in the brain, how it functions to alter our pain perceptions, why we have such a powerful mechanism, for example—they have said little more than "…and some other stuff happens in the head that explains everything else." As Wall himself remarks about the many areas implicated in inhibiting our sensations of pain, 'Unfortunately we know little of their relative importance and nothing of the actual circumstances in which they come into action.' (*The Myth of Pain*, 129.)

Hardcastle's own account, however, does not fare any better than gate-control theories with respect to any of the criticisms raised in this passage; a cursory glance through any standard textbook on pain will provide *vastly* more detail and complexity concerning gating mechanisms

---

[180] See, for example, Craig's model; e.g., Craig (2003a).
[181] She says she will offer three reasons (129), but I can only find two.
[182] Ibid., 129.

than anything Hardcastle provides about the PIS. Moreover, as discussed in chapter 3, Melzack and Wall bluntly acknowledge that they cannot explain all the inhibitory mechanisms relevant to nociceptive processing and explicitly acknowledge that the gating mechanism they identify is only one (among at least three) needed to explain all pains. Recognizing that there is still much to learn about how many inhibitory mechanisms there are and how they work, as Melzack and Wall do, appears warranted and appropriate. What is odd is that Hardcastle attacks Melzack and Wall on these grounds when her own model is at least as incomplete and promissory. Hardcastle's attack on gate-control theories on the grounds that they lack detail concerning inhibition, while positing an entire inhibitory system whose workings are at least equally vague, is unwarranted and inappropriate. She raises questions about the gating mechanism(s) that she never even attempts to answer of her own PIS. Why should this lack of detail be a damning objection to gate-control theories while leaving her own unscathed?

Hardcastle's attack here is all the more befuddling as she occasionally acknowledges herself how little we understand about the mechanisms that modulate nociceptive processing. As an example, consider the following:

> Though we only know a little about our pain inhibition in the brain (we do not know, for example, what counts as 'normal' function of our opioid analgesics, what environmental circumstances activate them, or how they actually affect pain sensations or behavior), consensus is emerging that we have several distinct inhibitory subsystem. (*The Myth of Pain*, 132.)

It is true that consensus is emerging that there are *several* neural pathways and mechanisms implicated in both inhibiting and amplifying nociceptive processing. Nonetheless, it remains unclear to *everyone* exactly how many such mechanisms and pathways there are or exactly how they work. Hardcastle is no exception—though she *is* exceptional in positing these still only vaguely identified mechanisms as a distinct system.

As mentioned, Hardcastle's second objection to gate-control theories is that they characterize pain experiences too subjectively. She writes:

> Second, and more important, even if we could get the details on some sort of central gating mechanism, this would not mean that pains are not largely subjective (which is but a step away from being purely subjective). If top-down cortical processes (which is what I take "purely subjective" to mean here) are mainly responsible for our sensations of pain, then the IASP subcommittee would be right and pains would be deeply peculiar. They would be the only perceptual experiences we have that are normal (that is, normal functioning creatures have them), natural (that is, not the product of tweaking something in a laboratory, but occurring in the wild, as it were), commonplace (we have them all the time), believed to be giving us information about the external environment (external to the brain, that is), but in fact do not. (*The Myth of Pain*, 129-130.)

Gate-control theory is here objected to for emphasizing top-down cortical processing. No reason is offered here for the puzzling claim that downstream cortical processing is problematically subjective. And, of course, even if this was somehow argued to make pain *largely* subjective, there is no reason to take the "step away" to *purely* subjective. So too, recall from chapter 3 that Melzack and Wall explicitly *reject* the IASP definition of pain. Most central for evaluation the attack: Activation in *Hardcastle's* offered model of the PIS *begins* with downstream cortical projections[183] —and it is the activities of the PIS that she appeals to to explain all of pain's oddities! Why does downstream cortical processing render gate-control theory problematically subjective, but leave her own unscathed? Again her objections to gate-control theory apply to her own model just as forcefully.

Hardcastle concludes her critique of gate-control theories noting that she "…advocate[s] dividing what Melzack and others have lumped into one cortically driven system into two separate systems: a nociceptor-driven pain sensory system (PSS) and a largely top-down pain

[183] Ibid., 133.

inhibitory system (PIS)."[184] Hardcastle's objections to gate-control theory, however, do not establish the supremacy of her own view. If gate-control theory is vague and mysterious, so is Hardcastle's. Neither, moreover, is gate-control theory more "subjective" than Hardcastle's dual-systems model; if Hardcastle's PIS is not spooky because it relies on downstream cortical processing, gate-control theories are not either. There seems no good reason to posit Hardcastle's two independent systems as against the (complex) single system of gate-control theory.

Perhaps we should not yet dismiss Hardcastle's independently functioning PIS just yet, however, as her central consideration for its independence is teleological. She writes that "[m]ore important than being anatomically distinct and dissociable from nociception, the inhibitory streams are teleologically distinct from our PSS system as well."[185] The claim of distinct teleology is peculiar, however, since she functionally defines the PIS *solely* by its relation to the PSS. Consider Hardcastle's own summary of the distinct teleology of the systems:

> A PSS and PIS then serve two different goals: the PSS keeps us informed regarding the status of our bodies. It monitors our tissues to maintain their intactness whenever possible. In contrast, the PIS shuts down the PSS when flight or fleeing is immanent, and then enhances the PSS response in moments of calm. If our brains are geared for motor control, then the dual pain system makes good biological sense. (*The Myth of Pain*, 134.)[186]

It is difficult to see why we should accept the PIS as a separate, independently functioning system with a distinct teleology when its only function posited by Hardcastle is inhibiting or amplifying activity in the PSS. When is it *ever* functioning independently *or* for an independent purpose?

---

[184] Ibid., 130.
[185] Ibid., 132
[186] Besides being a just-so evolutionary story, this sketch is unsupported at best and inaccurate at worst. Inhibitory mechanisms rarely *shut down* the PSS, though nociceptive information is almost, if not always, being modulated by *multiple* mechanisms. For more on the modulation of nociceptive processing see chapters 6 and 7 of *The Pain Textbook.*

Recall, however, that Hardcastle's initial reason for positing the PIS was that its independent functioning could explain the oddities of pain—including, but not limited to, its poor correlation with nociception. If Hardcastle's dual-systems model can provide these explanations, then there may be good reason to accept her positive view after all.

Hardcastle chooses self-injurious behavior, or SIB, as her key example of a puzzling pain phenomenon that her dual-systems model can adequately explain. She claims that positing the PIS allows us to understand SIB as a rational response to pain: "SIB is simply a way to force our brains to initiate a pain-inhibiting sequence."[187] Notice, however, that the PSS as identified by Hardcastle is *not* antecedently active in SIB. As Hardcastle notes, "Though we find few commonalities among those who engage in SIB, there are at least two: self-injurers feel psychological distress prior to harming themselves, and they feel better afterward."[188] This description may well be accurate, but it does not fit Hardcastle's model. People do not engage in SIB to comfort themselves in response to *nociceptive processing*. On the contrary, they comfort themselves *by* kicking off their nociceptive processing—i.e., by damaging their tissues. The precipitating psychological distress is undoubtedly physically realized in the self-injurers brain or nervous system, but it is *not* realized, much less initiated, by *nociceptive* processing.

One might respond by saying that the precipitating distress *caused* activity in the PSS, thereby compelling the self-injurer to initiate activity in the PIS to modulate it. Besides being antecedently implausible, this move is not open to Hardcastle. She earlier argues at some length that "…there is no good reason to hypothesize that any psychological factor or trait causes pain."[189] Hardcastle identifies a main goal of her book to be to establish that there is *no such*

---

[187] Ibid., 141.
[188] Ibid.,141.
[189] Ibid., 31.

thing as pain of psychogenic origin; real pain is always a function of activity in the PSS and the PIS, and never has a psychological origin. By Hardcastle's hypothesis, nothing psychological could *ever* initiate a pain experience. It is not clear, then, why Hardcastle consider people who engage in SIB as being in real pain at all, much less as the paradigm for a type of puzzling pain that her dual-system model explains. There is no way open to Hardcastle of explaining activating the PIS to quelch over activity in the PSS as a rational response since, prior to the SIB itself, there is no reason to think there *was* any activity in the PSS.

Going further, consider the evidence required to support the model: We would need evidence both that pain without tissue damage is explained by an underactive PIS relative to the PSS, and that tissue damage without pain is explained by an overactive PIS relative to the PSS. Consider phantom limb pain. Is the lack of correlation between tissue damage and the experience of pain in a non-existent leg somehow explained by under-activity of the PIS relative to activity in the PSS? Besides being at odds with leading theories of phantom limb pain, this is implausible; the relevant primary afferents are missing. For another case, consider headaches. Is the lack of correlation between tissue damage and the experience of pain explained by under-activity of the PIS? How? Hardcastle characterizes PIS activity solely as a response to activity in the PSS, but many of the puzzling cases are puzzling precisely *because* there is pain without PSS activity.

In sum, none of the reasons Hardcastle offers—anatomic, dissociative, teleological or explanatory—are compelling reasons to think of the PIS as an independent system. So too, while inhibitory mechanisms are crucial to our understanding of nociceptive processing, Hardcastle provides no compelling reasons for thinking that her dual-systems model fares better than the gate-control theories that also include these. Her objections to gate-control theory apply equally,

if not more forcefully, to her own model. All this may be overlooked if Hardcastle provided evidence that her dual-systems theory can make explanatory advances for pains that are puzzling from a perceptual-theory perspective, but her central case of SIB is unsuccessful and no other compelling cases are offered. It is, moreover, hard to see how successful explanations of these cases *could* be offered using the proposed model. I conclude that Hardcastle's dual-system model is undermotivated.

Having evaluated her positive proposal, I turn to Hardcastle's argument that the folk notion should be eliminated in favor of it. The first and most immediate problem with Hardcastle's suggestion is that she inconsistently characterizes the notion she wants eliminated. As Timothy Bayne points out in his review:

> According to Hardcastle, science tells us that pain sensations themselves have different dimensions, whereas the folk notion of pain requires that a pain is "… a single simple sense datum". (154) But Hardcastle does not characterize the folk-notion of pain sensations consistently, for she also says that "our commonsense views tie the different aspects of pain sensations into single *complex* percepts" (p. 151, my emphasis). (*Review—The Myth of Pain.*)

Bayne's objection here is on target. Which notion is the candidate for elimination?

Hardcastle's arguments for eliminativism of the folk notion of pain all rely on characterizing it as a *simple sensation* that is necessarily conscious and necessarily awful. How do her eliminativist arguments fare if we ignore her contradictory characterizations of the folk notion and focus on only the problematic notion that she thinks supports eliminativism?

I have repeatedly argued that I do not think the folk notion of pain is appropriately characterized as of a simple sensation that is necessarily conscious or necessarily awful. Hardcastle never gives any arguments for this problematic construal. As presented in chapter 2, the folk sometimes attribute pains to themselves and others in cases where the pains are not awful. Additionally, as Bayne points out in his review, the fluid folk use of the McGill Pain

Questionnaire in diagnostic practices also supports a folk understanding of pain that is complex and multidimensional. There seems no reason to characterize the folk as Hardcastle does for her eliminativist arguments.

Even if we *accept* this problematic characterization of the folk notion, Hardcastle's arguments for elimination are not successful. The first and most immediate problem here is another tension in the text, again pointed out by Bayne. In offering a positive proposal, Hardcastle does not appear to accept her own eliminativist conclusion:

> Take the *nature* of Hardcastle's elimination. On the one hand, she claims that "our folk-theory of pain is woefully inadequate. We might play at revising this theory, but any change that remains faithful to what we know about pain processing is gong to entail that pain [actually: 'pain'] no longer refers to a simple conscious percept. This sort of change crosses the line from mere revision to outright replacement" (p. 159). Yet Hardcastle remains quite happy to talk about pain and pain sensations. Eliminativism about *X*s is normally taken to be the claim that there are no *X*s, but Hardcastle seems to think that there are pains. Indeed, in her final chapter she argues that we should do more to treat the pain of children and infants. So the sense in which she is an eliminativist eludes me. (*Review—The Myth of Pain.*)

I think Bayne's criticism again hits the mark. Hardcastle says that we should eliminate the folk notion of pain. If so, however, then why does she bandy the term 'pain' about with such frequency and what is its referent when she does? If pains do not exist, then what is her dual-systems model a dual-systems model *of*?

I take it that Hardcastle's reply should be that her proposed referent of 'pain' is different than the referent of 'pain' in folk theory. The folk referent does not exist and, moreover, the theory by which the term's referent is derived is bad. The new proposed referent, however, is scientifically respectable and derived from a good theory. By using the folk notion of pain (by hypothesis, a simple sensation that is necessarily conscious and necessarily awful) we discovered that pain is *actually* multidimensional. The conscious, awful, sensation the folk think is pain is

only one aspect of a complex process. Since we used the initial folk term 'pain' in the process of identifying the new referent, we can go on using that *term,* but we have actually changed the subject. This is my best attempt to rescue Hardcastle's folk eliminativism from Bayne's damning criticism, but it does not work.

Why should we eliminate instead of revising the folk notion of pain? Put more fatally: Isn't Hardcastle's own positive proposal revisionary instead of eliminativist? Hardcastle offers a positive account of the nature of *pain.* If 'pain' can be salvaged for science, then why can't it be salvaged for everyday purposes? When Hardcastle addresses this objection directly, she construes it this way:

> Defenders of a revised folk psychology answer that … the new term and concomitant theory would be a mere revision of 'pain.' Without better exposition of what the replacement for our folk psychology term will be (and how it will be radically different), we simply cannot tell what the future holds for our commonsense theory of pain: simple revision, peaceful coexistence, or outright replacement. (*The Myth of Pain*, 159.)

Why is her proposal the eliminativist one she claims it to be instead of a revisionary proposal? Well:

> The solution for the eliminativist, of course, is to offer a serious replacement for our folk terms. Chapters 5 and 6 outline a substantive, biologically based theory of pain processing. The combination of a PSS and a PIS account for the appearance of pain sensations being disconnected from physiology. …
> Though other aspects of our folk psychology may not be wrong, our folk theory of pain is woefully inadequate. We might play at revising this theory, but any change that remains faithful to what we know about pain processing is going to entail that pain no longer refers to a simple conscious percept. This sort of change crosses the line from mere revision to outright replacement. (*The Myth of Pain*, 159.)

This response will not do. First, no reason is offered for this last sentence. Hardcastle is right that "…a good theory of pain that remains faithful to what we know about pain processing" will not be a theory of pain as of a simple, conscious percept. (On my view, that is not the folk theory

either.) Why, however, does accommodating that truth entail *replacing* the folk notion of pain instead of merely revising it? Hardcastle merely stipulates that accommodating this change would be eliminativism, in the very context of the objection that it is only a revision.

The above response to the revisionary challenge, moreover, leaves the question of what her positive theory is a theory *of* given her eliminativism unanswered. Here Bayne's objection surfaces most damnably. Hardcastle maintains that the folk concept of pain has no referent. Thus she writes:

> …if what we refer to as pain turns out to be neither a conscious experience nor a perceptual simple, then what we are trying to point to in the world with the word does not exist, for these are the two defining features of our folk notion of pain. (*The Myth of Pain*, 154.)

If so, however, then the scientific research she appeals to is not *about* whatever the folk (think they) are referring to. One does not argue that there is no such thing as phlogiston by offering a scientifically respectable theory of phlogiston! She cannot consistently maintain that pain does not exist, while still maintaining that *pain* has been *discovered* to be anything. Bayne's objection is fatal. Hardcastle's two central claims contradict each other; her eliminativism and positive proposal cannot both be true. Either she is actually advocating only revision and not elimination, or she *is* advocating eliminativism and her dual-systems theory is not a theory of pain.

Given my arguments against her positive proposal, I will consider the arguments for eliminativism independently. If we ignore her contradictory positive proposal, how do her eliminativist arguments fare?

Hardcastle offers two central reasons for elimination of the folk notion, and the second is

that our language for talking about pain is inadequate.[190] She writes:

> There is a second reason for redesigning our ways of understanding our experiences of pain. The language we have for expressing our propositions concerning pain sensations is very crude. … We simply cannot express in a clear and unambiguous fashion how pain, in all its complexity, feels. (*The Myth of Pain*, 151.)[191]

And again:

> We simply do not have the language to express all the dimensions of our pain experiences. The descriptors are either metaphorical or nonexistent. Our folk theory of pain needs to be replaced by something commensurate with the phenomenology. (*The Myth of Pain*, 153.)

This reason is puzzling: How does the paucity of our language to describe our sensations support that 'pain' should be eliminated or that pains as the folk conceive of them do not exist? It is hard to see.

How does denying that pains exist help us better express them? Is it supposed that the term 'pain' gets in the way of our expressing our pain experiences? On the face of it, that is implausible. Is it supposed that if our folk notion of pain was worth retaining, then we would already be better able to describe our pain experiences? Why? Perhaps the reason is the suggested one that replacing our folk notion of pain with Hardcastle's proposed dual-systems model will somehow provide new and more adequate descriptors. Beyond the contradictions mentioned above that this will entangle us in, it is hard to see what *new* or *better* descriptors would be gained by elimination. Hardcastle does not provide a schema for describing our pains in a way that would be more "commensurate with the phenomenology." Most important, novel and superior descriptors would not support abandoning the old terms instead of adding to them.

---

[190] Hardcastle also argues that our folk notion of pain is essentially, though tacitly, dualistic. She merely stipulates this, and even if she is correct, here again Hardcastle gives no reason for why this should prompt elimination instead of mere revision.

[191] 'Redesigning' here sounds more like revision than elimination, but she is explicit that she is here offering a reason for elimination.

Hardcastle herself retains the term 'pain'! Even if our ability to express our pains is as impoverished as Hardcastle contends, it is unclear why eliminativism would help.

Her first offered reason for folk eliminativism is more compelling, however, and usefully illuminates the contrast between Hardcastle's position and my own. According to Hardcastle, pain research shows that the folk notion of pain conflates dimensions of our pain experiences that should be kept distinct:

> Some may argue that pain then is just a simple term for a complex sensation. That is, we have just uncovered different aspects to a single perception, instead of entirely different experiences. Hence, it would make good folk and scientific sense to keep one name for the complex percept.
> I disagree. The four dimensions of our pain sensations [the bodily location of the sensation, its somatic quality, the feeling of suffering, and negative reactions to the feeling of suffering] should be kept distinct. They are not like the color and shape of an object or the pitch and loudness of a sound. Our perceptions of objects and sounds are comprised of complex units. We cannot see a shape without also seeing that shape as some color. We cannot hear a loud noise without also hearing that noise at some pitch. If a brain is not given some of the information regarding the structure of an object or a sound, then it simply fills in what is missing.
> In distinction, pain sensations function quite differently. We quite often get one aspect of pain without the others. Moreover, if we are lacking an aspect… then our brains feel no obligation to fill in the missing pieces. Though our commonsense views tie the different aspects of pain sensations into single complex percepts, our brains do not (*The Myth of Pain*, 151.)

While I reject the claim that the folk notion of pain is as of a simple sensation, I have argued in the previous chapters that the folk notion of pain is not useful for science. Let us agree with Hardcastle that pain science has illuminated that the referent of the folk notion does not pick out anything useful for scientific theorizing about pain—and bracket that this contradicts Hardcastle's positive proposal. Why does Hardcastle think this presses elimination of the *folk* notion?

Hardcastle seems to think that if something is (or is not reducible to something) appealed to in scientific laws, then it does not exist. Accordingly, about semantic content she writes:

> It is entirely possible that science ultimately will not use the notion of semantic content in any of its laws. If this were to occur, we would have to conclude that representations are not scientific entities and any explanations of behavior that rely on the intentional in the explanans will have to be dismissed as incorrect just as we now dismiss explanations that use witches or phlogiston. (*The Myth of Pain,* 55.)

If semantic content is not employed in any of science's laws, then it will go the way of witches or phlogiston.

The methodological plot thickens, however, when she claims that even if something does not exist, reference to it might still be acceptable for folk purposes. Thus she writes:

> So far as I can tell, content not being a scientific term will have little impact on our common sense explanations of one another. 'Folk' psychology, like other folk 'sciences,' need not be concerned with facts of the matter. It could function as well (or as poorly) as it ever did, regardless of whether intentional states really exist. (*The Myth of Pain*, 56.)

The overall methodology embodied in these passages and others seems to be that whether something exists depends on whether it is a natural kind; i.e., whether it is usefully referred to in science's generalizations/used in science's laws.[192] Whether something exists, however, is held to have little impact on our folk psychology, because folk psychology "need not be concerned with facts of the matter."

Bracketing some very dubious assumptions embedded here (what could it mean to say that folk, or any other, theory is not concerned with facts of the matter?), if this is Hardcastle's methodological approach to the relationship between eliminativism of folk notions and natural kinds, why is pain supposed to go the way of the dodo? Why should we eliminate the folk notion of pain even *if* 'pain' is not a "scientific term"?

This question highlights the differences between Hardcastle's view and my own, because

[192] See again chapter 1 for my preferred characterization of natural kinds and their relationship to science's generalizations, laws, and models.

while I *do* think that pain is not a natural kind, I also think that we should nonetheless *retain* the folk notion of it for everyday purposes. This position seems open in Hardcastle's general methodology (though, again, based on some dubious assumptions about folk psychology). In the case of pain, however, she directly addresses and rejects it.

She presents this contrasting view as an objection to her folk eliminativism in the following passage:

> …detractors point out that even if a completed psychology does not rely on the propositional attitudes or consciousness or some other defining characteristic of our folk theories, that would not entail that those sorts of mental states do not exist; they would just not be referred to in scientific discourse. … We might not have a mental science or laws about pain, but our folk psychology could still be used as it is now, in our everyday explanations of behavior. It works well enough, in our specific cultural contexts, to get the job done. Hence, there is no need for elimination or replacement. (*The Myth of Pain*, 156.)

As it is not, to my knowledge, elsewhere published in the literature, I was grateful to have come across a characterization of my own position being taken seriously. Hardcastle's response to it is unpersuasive:

> There is something undoubtedly right about this charge. …not enough to show that our folk ways of understanding pain are adequate. … varied expressions for varied phenomenology does not guarantee that the descriptors are sufficient within their contexts. I argue that ours, anyway, are not, for at least two reasons. First, people with pain experiences that do not fit under our folk theories are assumed to have some psychiatric problem. Chronic pain without a diagnosed physical cause is deemed a mental disorder. Nonpainful self-injuries are considered a mark of Borderline Personality Disorder. Second, the fact that different chemicals affect the qualitative experience of pain differently cannot be expressed using our folk measure of pain intensity, namely, how much it hurts. This has had serious consequences for patients trying to relieve their pain or suffering with drugs. I claim that there is a fundamental inadequacy to our folk theories. The burden is now on the folk psychologist to demonstrate how—despite appearances—our folk theories of pain are actually successful. (*The Myth of Pain*, 157.)

The response to the objection that we ought to eliminate 'pain' from scientific theorizing but not from everyday life, then, is that our folk understanding of pain is inadequate for (at least) two

reasons: improper diagnosis and the inadequacy of intensity metrics to capture the complexity of pain phenomenology for treatment purposes. Both of these reasons, however, confuse the adequacy of the folk theory with that of the scientific theory.

Hardcastle blames the purported inadequacies of the folk notion for the inadequacies of diagnosis and treatment, but the identified inadequacies are not even *part* of folk theory. Why blame the folk notion for inadequacies in diagnosing chronic pain? The folk do not diagnose each other with chronic pain or borderline personality disorder. The folk notion, as argued in chapter 2, countenances chronic pain as *real pain* despite the lack of (known) tissue damage. It is Hardcastle's own positive account, and others like it, that illegitimately try to shoehorn pain into a unified category and thereby wreak havoc for our understanding of chronic pain—however sincere her intentions to the contrary. The folk notion of pain is not intended for diagnosis and it is inappropriate to eliminate that notion because it fails to adequately provide it.

Similar comments apply to the problematically simple scales of pain intensity. These scales are not *folk* products, so why should the folk be blamed for their inadequacy? The folk describe their pain using a rich variety of descriptors. If the doctor's metrics are overly simplistic, how could the folk theory be to blame? When speaking to each other in everyday life the folk do not limit themselves to expressing and reporting their pains on a simple intensity scale of how much it hurts from 1 to 10. If they are allowed, they also do not so restrict themselves when expressing and reporting their pains to their doctors. Moreover, simple intensity metrics are not the most prevalent or respected tool for pain diagnostics: the McGill Pain Questionnaire is. The questionnaire, as discussed in previous chapters, prompts patients to describe their pain using a number of descriptors for the multiple dimensions of pain experiences. Of course, sometimes doctors *do* employ simple intensity metrics. While it may sometimes be appropriate to do so (an

issue of no immediate concern here), there is no reason to think the folk even *approve* of these metrics.

As a brief detour, consider the following two anecdotes of the folk *mocking* simple intensity metrics. First, an xkcd comic:[193]

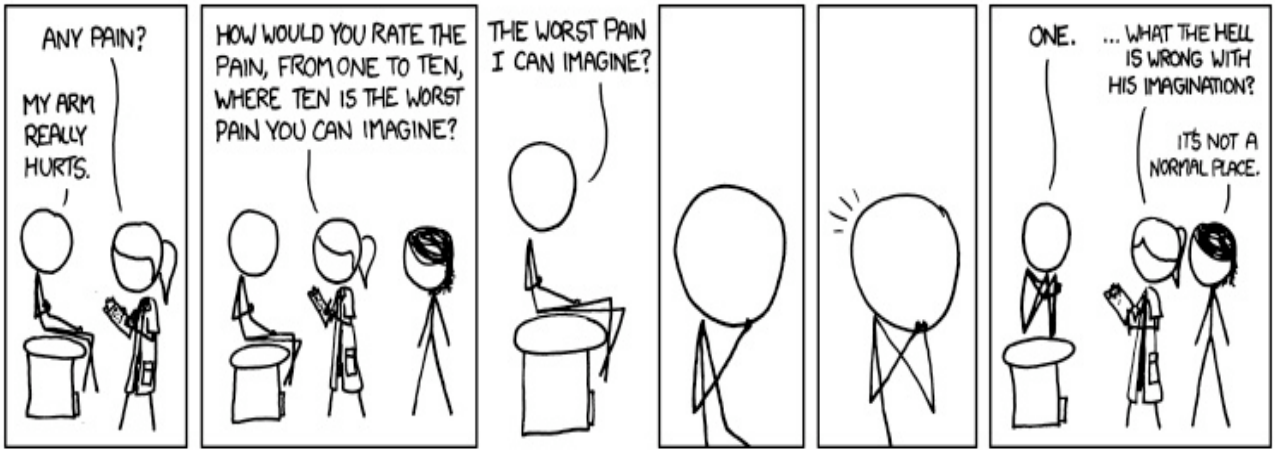

Figure 6.1: Imagining Pain. Taken from xkcd.com.

Second, part of a blog post[194]:

[193] Thanks to David Pereplyotchik for bringing this to my attention. This comic and others can be found at: http://xkcd.com/883/

[194] Thanks to Aaron Holt for bringing this to my attention. The entire blog post is available at: http://hyperboleandahalf.blogspot.com/2010/02/boyfriend-doesnt-have-ebola-probably.html.

You've probably seen some version of that chart before. You may also have noticed how inadequate it is at helping you. Based on the faces, this is my interpretation of the chart:

0: Haha! I'm not wearing any pants!

2: *Awesome*! Someone just offered me a free hot dog!

4: Huh. I never knew that about giraffes.

6: I'm sorry about your cat, but can we talk about something else now? I'm bored.

8: The ice cream I bought barely has any cookie dough chunks in it. This is not what I expected and I am disappointed.

10: You hurt my feelings and now I'm crying!

Figure 6.2: Intensity Chart Interpretation. Taken from hyperboleandahalf.blogspot.com.

The blog's author goes on to comment that "None of that is medically useful…" and creates their own intensity metric explicitly (and hilariously) mocking the doctor's overly simplistic one. Commenters on the post agree that the simple intensity metrics are ridiculous and unhelpful; e.g., "…everytime I go to the hospital (read: a lot) they ask me to rate the pain. And I'm all, well my 8 may not be your 8 so can we discuss this further? Then they roll their eyes…," or "I have *multiple* chronic pain issues, and have always hated these stupid charts," and my favorite, "I have never in my life seen one of those charts before, they seem kinda rediculous, can't people

just explain using their like mouths and stuff?"

These comedic anecdotes are obviously comedic and anecdotal, but I think they usefully illustrate that there is no reason to think that the *folk* are to be blamed for the overly simple intensity metrics. If these anecdotes are indicative of the folk perspective more broadly, these simple intensity metrics are a joke. To impugn the folk notion of pain for the inadequacies of simple intensity metrics is itself a joke.

After identifying the supposed inadequacies of the folk theory for diagnosis and treatment, Hardcastle asserts that she has at least shifted the burden of proof to the supporter of the utility of the folk notion of pain. If my refutations above are successful, then she has not: her offered reasons may provide further support for the claim that scientific models of pain are inadequate, but provide no support for her folk eliminativism.

In sum, the position developed in this dissertation, is the opposite of Hardcastle's: our folk notion of pain should not be eliminated and there is no useful notion of pain for scientific theorizing. I agree with Hardcastle that pain *is* a complex, multidimensional experience and that ignorance of this has led many philosophers to illegitimately characterize pain as some one dimension of its complexity. Unfortunately, however, Hardcastle's own proposal commits a variant of this same mistake: the wide range of pain experiences are *not* all explained by interaction of the PSS and the PIS. There is, moreover, inadequate evidence that we even *have* a PSS (a *pain* sensory system as against pathways that process nociceptive information) or an independent PIS (as against numerous mechanisms that modulate nociceptive processing). Hardcastle's arguments that we should eliminate the folk notion of pain are similarly weak. Putting aside the contradiction between her eliminativist and positive proposals, the reasons offered for folk elimination support, if anything, our *scientific* theories of pain. If our language

for expressing our pains is inadequate, denying that pain exists, or ridding us of the term 'pain' will not help. If our diagnosis and treatment of pain conditions are inadequate, the finger should be pointed squarely at the scientific theories.

Hardcastle has done the field a service by nudging debates about the nature of pain further from the armchair and deeper into the messy world of pain research, and that is laudable. Based on this research, however, Hardcastle concludes that there is a useful notion of pain for scientific purposes and that the folk notion should be eliminated. The appropriate, however, are the opposite: we should eliminate the term from scientific theorizing while retaining it for folk purposes.

4. Conclusion

In this chapter, I have argued that the prominent arguments folk eliminativism offered by Dennett and Hardcastle are unsuccessful, but my own argument against pain's naturalness may be taken as an argument for folk eliminativism. If pain is not a natural kind, then it is not a suitable *folk* kind; if it is not a useful referent for science, then it does not exist. On the contrary, I maintain that some non-natural kinds exist and that a term that is useless for science may nonetheless be useful for everyday life. So it is with pain, 'pain,' and terms referring to pain "types."

In chapter 5, I argued that no useful generalizations would be lost by ridding our scientific inquiry of 'pain' and terms for pain "types"—all candidate generalizations would be more reliable and successful if we employed only terms referring to the multiple mechanisms implicated in pain experiences. If I am right that neither pain nor pain "types" are usefully referred to in scientific inquiry, however, then why keep referring to them in daily life? Should we start referring only to the mechanisms implicated in pain experiences in daily life as well as

in scientific generalizations?

In general, we should hang on to our folk notions unless we have good reason to abandon them—non-naturalness is not a good enough reason. In particular, 'pain' has a useful folk referent—it refers to those experiences that paradigmatically include characteristic sensation, perception, emotion and avoidance behaviors. That this complex admits of no mechanistic explanation does not 1undermine its utility for folk purposes—this despite my wholehearted agreement that its lack of a mechanistic explanation *does* undermine its utility for scientific generalizations.

This need not trouble us if we attend to the fact that the purposes of the folk are not the purposes of the scientist. There is a folk notion of pain. Whether that notion is useful for *folk* purposes is not appropriately evaluated based on its utility for *scientific* purposes. When I say I am in pain, I tell you something. Indeed, I can *lie* to you and say that I am in pain when I'm not. Why should these mundane claims be denied merely because there is no mechanistic explanation of pain? None of the prominent reasons offered for eliminating the folk notion of 'pain' are successful. We can go further: What *is* the utility of the folk notion of pain?

The folk notion of pain is, in short, useful for communicating with other folk. It allows us to become aware of, self-ascribe, and communicate to others the bad things happening to us presently—both physical and emotional. And that is enormously useful. It is instrumental in our search for medical attention, compassion, and sympathy. It facilitates our ability to protect injured limbs, to limit how far we push our bodies, and to identify potentially toxic chemicals and stimulation. It is, I suspect, the very promiscuity of the concept that allows it so *much* utility in our everyday interactions and in our self-monitoring. The folk notion of pain, as Dennett aptly pointed out, has much slack. We usefully speak of emotional pain, physical pain, stinging or

throbbing pain, growing pains, and even *fiscal* pain—even pains that do not bother us. The folk descriptions of pain do not always supply categories for medical diagnosis, but they are always quick and dirty, and thereby *useful* for getting us to a doctor and getting a doctor's attention. Given all the *prima facie* uses of our folk notion of pain we ought to continue to embrace it until we have good reason to reject it. Pains' non-naturalness is not a good enough reason. as it does nothing to undermine this utility.

Pain or pain "types" lack a mechanistic explanation because each token pain is explained by the *idiosyncratic* convergence of *multiple* mechanisms. It is, ultimately, the idiosyncrasy that undermines the success of any generalizations, but this idiosyncrasy does not undermine the folk purposes identified above. That *my* pain *right now* has an idiosyncratic explanation does not undermine that I have usefully communicated something to you about my current situation by referring to it.

Ontological worries may nonetheless loom large. If a kind is non-natural then it may further be claimed that its members therefore do not exist. Dennett, as discussed in section 2, makes a similar move: no mechanical explanation, then no existence. I see *no* reason to accept this inference, however, and many examples support rejecting it.

Consider first the days of the week. Tuesday, I take it, is not a natural kind by the scientific utility criterion developed in chapter 1. Tuesdays are not usefully referred to in the generalizations of any relevant science—there is no usually co-occurrent clustering of Tuesday properties whose co-occurrence is explained by an underlying, homeostatic mechanism. If 'Tuesday' does not appear in any mature science, should we conclude that Tuesdays do not exist? Absurd. There are, of course, many ontological issues in this country far beyond the province of this dissertation. For purposes here: I am happy to conclude that the ontological

status of pains is as robust as that of Tuesdays. One may say that there are not *really* any Tuesdays if they like, but Tuesdays are real enough for me—even if they are never usefully referred to in the generalizations of any relevant science. Denying that there are *any* non-natural kinds is unwarranted and extreme.

While examples like the days of the week counter the general inference from non-naturalness to non-existence, it must be admitted that Tuesdays and pains are arguably different in germane. Pains and pain "types" are folk-psychological kinds that are arguably best characterized *functionally*—that is, *causally*. While days of the week are useful for organizing behavior, they are not characterized in a functional, causal way. It is the *causal* characterization of folk-psychological kinds that may lead one to claim that a folk-psychological kind that is not natural does not exist at all. This too, I think, is too extreme. There is no good reason to deny that there are any non-natural folk-psychological kinds.

Consider friends. Friends are a social, but (at least arguably also) folk-psychological kind that provides are closely parallel to pan and pain "types" as I have argued them to be. Friends are a useful category for everyday purposes. When I call someone my friend I say something meaningful, and I can lie and say that someone is my friend when they are not. A cluster of usually co-occurrent properties no one of which is necessary or sufficient characterizes friends. As I have argued is true of pain, the usual co-occurrence of the properties of friends are presumably not, however, explained by any one, underlying homeostatic mechanism. No one thing either reliably causes or destroys friendships. Instead, whether someone is a friend at least seems to be explained by an *idiosyncratic* combination of *multiple* factors. As such, scientific generalize about friends would be rather poor. Nonetheless, of course, friendships not only exist but are a real and important phenomenon in our lives. Their idiosyncrasy blocks useful scientific

generalizations about them and they do not admit of a mechanical explanation. The idiosyncrasy of friendships arguably, however, also makes them all the more valuable.

Let us review. In this dissertation, I argued that neither pain nor any pain "type" is a natural kind. In chapter 1, I argued that a kind is natural insofar as it is usefully referred to in the generalizations of the relevant science. Following Boyd, I took the best indicator that a kind *is* usefully referred to by science's generalizations to be that it is explicable, at least in principle, by an underlying, homeostatic mechanism. I then claimed that pain and pain "types" fail to be natural using this criterion because accounts of pain inevitably either under-specify the phenomenon, or lump many phenomena with different mechanistic explanations together. Chapter 2 argued that unitary accounts of pain commit the first mistake and chapters 3 and 4 argued that complex accounts commit the second. In chapter 5, I denied that the arguments in chapters 3 and 4 supported only a pluralism about pain; the mechanical differences across token pain experiences really do undermine useful references to pain or pain "types" in any scientific generalizations. In the current chapter, I argued that despite the non-naturalness of pain and pain "types," pain is nonetheless very real and references to it and its types should be retained for everyday purposes; folk eliminativism about pain should be rejected.

The dissertation as a whole, then, constitutes a case study of an ontologically robust folk-psychological kind that scientific inquiry has revealed is not natural. Pain and pain "types" are important, yea central, for daily life, but they are not natural kinds.

In the current philosophical climate, the very possibility of such a kind may seem incredible. In *Empiricism and the Philosophy of Mind*, Sellars famously claimed that "…in the dimension of describing and explaining the world, science is the measure of all things, of what is that it is, and of what is not that it is not." In this spirit, one may maintain that *something* must

have gone wrong in this dissertation. Insofar as pain and pain "types" *do* exist, then as good naturalists we have every reason to infer that science can offer fruitful generalizations about them—even if it has not yet managed to do so. Fully assuaging this worry requires settling issues at the heart of philosophy of mind, language *and* science—a task far greater than can be accomplished here. For now, I will simply say that I think it is time to re-think some of the assumptions embedded in this worry in the light of the arguments I have presented.

I argue that each token pain experience is explained by the idiosyncratic convergence of the activity of multiple mechanisms, such that useful scientific generalizations about pain or pain "types" are blocked. This is an instance of a general truth: Some things are idiosyncratic and individual. We need a philosophical methodology that can not only recognize their existence, but that can speak insightfully about them. Without that (to reference another Sellarsian insight), our attempts to harmonize the manifest and scientific images will remain unsuccessful.

Bibliography

Al-Chaer, Elie D. and Richard J. Traub, "Biological basis of visceral pain: recent developments," *Pain,* 96 (2002): 221-225.

Allen, Garland E., "Mechanism, vitalism and organicism in late nineteenth and twentieth-century biology: the importance of historical context," *Studies in History and Philosophy of Biological and Biomedical Sciences*, 36 (2005): 261 283.

Anand, K.J. and KD Craig, "New Perspectives on the Definition of Pain." *Pain* 67 (1996): 3-6.

Anand, K.J.S., and Kenneth D. Craig, "Re: Reply to Letters to the Editor from Merskey & Wall," *Pain* 67 (1996): 210.

Anderson, D. Bryant and James W. Pennebaker. "Pain and pleasure: alternative interpretations of identical stimulation," *European Journal of Social Psychology*, 10 (1980): 207-212.

Armstrong, D.M. *Bodily Sensations.* New York: Humanities Press, 1962.

Armstrong, D.M. "Is Introspective Knowledge Incorrigible?" *Philosophical Review,* 72 (1963): 417-432.

Armstrong, D.M. and Norman Malcolm. *Consciousness and Causality: A Debate on the Nature of Mind*. Oxford: Blackwell, 1984.

Aristotle. *Metaphysics* edited and translated by John Warrington. London: Dent; New York: Dutton, 1956.

Aydede, Murat, "Review of *Feeling Pain and Being in Pain,* by Nikola Grahek." *Notre Dame Philosophical* Reviews, 2004.

Aydede, Murat, "Is Feeling Pain the Perception of Something?" *The Journal of Philosophy,* 106 (2009): 531-567).

Aydede, Murat. "Introduction: A Critical and Quasi-Historical Essay on Theories of Pain." In *Pain: New Essays on Its Nature and the Methodology of Its Study*, edited by Murat Aydede, 1-58. Cambridge, MA: MIT Press, 2005.

Aydede, Murat and Donald D. Price. "Introspection and Unrevisability: Reply to Commentaries," In *Pain: New Essays on Its Nature and the Methodology of Its Study*, edited by Murat Aydede, 315-324. Cambridge, MA: MIT Press, 2005.

Aydede, Murat, "Pain, Philosophical Aspects of." In *Oxford Companion to*

Consciousness, edited by Tim Bayne, Axel Cleermans and Patrick Wilken. Oxford: Oxford University Press, (forthcoming).

Aydede, Murat and Guven Guzeldere, "Some Foundational Problems in the Scientific Study of Pain," *Philosophy of Science*, 69 (2002): 265-283.

Baker, Jason M., "Adaptive speciation: the role of natural selection in mechanisms of geographic and non-geographic speciation," *Studies in History and Philosophy of Biological and Biomedical Sciences*, 36 (2005): 303-326.

Barber, Theodore Xenophon. "The Effects of "Hypnosis" on Pain: A Critical Review of Experimental and Clinical Findings," *Psychosomatic Medicine,* 25, no. 4 (1963): 303-333.

Barden, Jodie, Jayne E. Edwards, Henry J. McQuay and R. Andrew Moore, "Oral valdecoxib and injected parecoxib for acute postoperative pain: a quantitative systematic review," *BMC Anesthesiology*, 3 (2003): 1-9.

Baumgärtner, Ulf, Walter Magerl, Thomas Klein, Hanns Christian Hopf, and Rolf Detlef Treede, "Neurogenic hyperalgesia versus painful hypoalgesia: two distinct mechanisms of neuropathic pain," *Pain* (2002): 141-151.

Bechtel, William and R.C. Richardson, *Discovering Complexity: Decomposition and localization as strategies in scientific research.* Princeton: Princeton University Press, 1993.

Bechtel, William and Adele Abrahamsen, "Explanation: a mechanist alternative," *Studies in History and Philosophy of Biological and Biomedical Sciences*, 36 (2005): 421-441.

Bechtel, William. *Mental Mechanisms: Philosophical Perspectives on Cognitive Neuroscience.* New York: Taylor and Francis Group, 2008.

Berthier, Marchelo, Sergio Starkstein, and Ramon Leiguarda, "Asymbolia for pain: A sensory-limbic disconnection syndrome, "*Annals of Neurology*, 24 (1988): 41-49.

Block, Ned, "On a Confusion About a Function of Consciousness," *Behavioral and Brain Sciences*, 18 (1995): 227-287.

Block, Ned, "Consciousness, accessibility, and the mesh between psychology and neuroscience," *Behavioral and Brain Sciences,* 30 (2007): 481-548.

Boyd, Richard, "Realism, Anti-Foundationalism and the Enthusiasm for Natural Kinds," *Philosophical Studies* 61 (1991): 127-148.

Bogen, Jim, “Regularities and causality; generalizations and causal explanations,” *Studies in History and Philosophy of Biological and Biomedical Sciences*, 36 (2005): 397-420.

Boyd, Richard, “Kinds, Complexity and Multiple Realization: Comments on Millikan’s “Historical Kinds and the Special Sciences”,” *Philosophical Studies*, 95 (1999): 67-98.

Brand, Paul and Philip Yancey. *The Gift of Pain*. Grand Rapids, MI: Zondervan Publishing House, 1997.

Burger, Jerry M., “Replicating Milgram: Would People Still Obey Today?” *American Psychologist,* 64 (2009): 1-11.

Carruthers, Peter, “Suffering without subjectivity,” *Philosophical Studies*, 121 (2004): 99-125.

Clark, Austen. *Sensory Qualities*. Oxford: Clarendon Press, 1993.

Clark, Austen, “Painfulness Is Not a Quale.” In *Pain: New Essays on Its Nature and the Methodology of Its Study*, edited by Murat Aydede, 177-197. Cambridge, MA: MIT Press, 2005.

Clore, et al. "Affective Feelings as Feedback: Some Cognitive Consequences." In *Theories of Mood and Cognition: A User's Guidebook,* edited by Leonard L. Martin and Gerald L. Clore, 27-62. Mahwah, New Jersey: Lawrence Erlbaum Associates, 2001.

Coburn, Robert C. "Pains and Space," *Journal of Philosophy,* 63, no.13 (1966): 381-396.

Coghill, Robert C. "Pain: Making the Private Experience Public." In *Pain: New Essays on Its Nature and the Methodology of Its Study*, edited by Murat Aydede, 299-306. Cambridge, MA: MIT Press, 2005.

Craig, A.D. "A new view of pain as a homeostatic emotion," *TRENDS in Neurosciences,* 26 no. 6 (2003a): 303-307

Craig, A.D. “Pain mechanisms: labeled lines versus convergence in central processing,” *Annual Review of Neuroscience,* 26 (2003b): 1-30.

Craig, A.D. “Interoception: the sense of the physiological condition of the body,” *Current Opinion in Neurobiology*, 13 (2003c): 500-505.

Craver, Carl F., “Beyond reduction: mechanisms, multifield integration and the unity of neuroscience,” *Studies in History and Philosophy of Biological and*

Biomedical Sciences, 36 (2005): 373-395.

Craver, Carl F. "When mechanistic models explain," *Synthese*, 153 (2006): 355-376.

Craver, Carl F.. *Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience.* New York: Oxford University Press, 2007.

D'Amico, Robert. "Sensations and Methodology," In *Pain: New Essays on Its Nature and the Methodology of Its Study*, edited by Mur2007at Aydede, 291-299. Cambridge, MA: MIT Press, 2005.

Dallenbach, K.M. "Pain: History and Present Status," *American Journal of Psychology*, 3 (1939): 331-347.

Darden, Lindley, "Relations among fields: Mendelian, cytological and molecular mechanisms," *Studies in History and Philosophy of Biological and Biomedical Sciences*, 36 (2005): 349-371.

Davis, Karen D., Chun L. Kwan, Adrian P. Crawley, and David J. Milkus, "Functional MRI Study of Thalamic and Cortical Activations Evoked by Cutaneous Heat, Cold, and Tactile Stimuli," *Journal of Neurophysiology*, 80 (1998): 1533-1546.

Dennett, D.C.. "Why You Can't Make a Computer that Feels Pain," In *Brainstorms: Philosophical Essays on Mind and Psychology*, edited by Daniel C. Dennett, 190-229. Cambridge, MA: MIT Press, 1981.

Dennett, D.C. "Personal and Sub-personal Levels of Explanation." In *Philosophy of Psychology: Contemporary Readings*, edited by Jose Luis Bermudez, 17-21. New York, NY: Routledge, 2006.

Dennett, D.C. *Consciousness Explained.* New York: Back Bay Books, 1991.

Des Chene, "Mechanisms of life in the seventeenth century: Borelli, Perrault, Regis," *Studies in History and Philosophy of Biological and Biomedical Sciences*, 36 (2005): 245-260.

Devitt, Michael. *Realism and Truth*. 2nd ed. Princeton, NJ: Princeton University Press, 1991.

Devitt, Michael. "Natural Kinds and Biological Realisms." Forthcoming in *Carving Nature at its Joints: Topics in Contemporary Philosophy, vol 8.*

Devitt, Michael. "Resurrecting Biological Essentialism," *Philosophy of Science*, 75 (2008): 344-82.

Dretske, Fred. "The Epistemology of Pain." In *Pain: New Essays on Its Nature and*

the Methodology of Its Study, edited by Murat Aydede, 59-73. Cambridge, MA: MIT Press, 2005.

Dupre, John. "Natural Kinds and Biological Taxa," *The Philosophical Review*, 90, no.1 (1981): 66-90.

Dupre, John. "Is 'Natural Kind' A Natural Kind Term?," *Monist*, 85 no. 1 (2002): 29-49.

Edwards, Jayne E., Anna D. Oldman, Lesley A. Smith, Dawn Carroll, Philip J. Wiffen, Henry J. McQuay and R. Andrew Moore, "Oral aspirin in postoperative pain: a quantitative systematic review," *Pain* 81 (1999): 289-297.

Fine, Kit. "Essence and Modality: The Second Philosophical Perspectives Lecture," *Philosophical Perspectives*, 8, (Logic and Language) (1994): 1-16.

Finnerup, Nanna B., and Troels S Jensen, "Mechanisms of Disease: mechanism based classification of neuropathic pain—a critical analysis," *Nature Clinical Practice: Neurology*, 2 no. 2 (2006): 107-115.

Fodor, Jerry A. "Special Sciences," In *RePresentations: Philosophical Essays on the Foundations,* edited by Jerry Fodor, 127-145. Cambridge, MA: MIT Press, 1981.

Fodor, Jerry A. "What is Universally Quantified And Necessary And A Posteriori And It Flies South In The Winter?," *Proceedings and Addresses of the American Philosophical Association,* 80, no. 2 (2006): 11-24.

Frances, Allen, and Leonard Gale. "The proprioceptive body image in self-object differentiation: A case of congenital indifference to pain and head-banging." *Psychoanalytic* Quarterly, 44, no. 1 (1975): 107-126.

Freeman, Walter and James W. Watts. *Psychosurgery in the Treatment of Mental Disorders and Intractable Pain*. Springfield, IL: C.C. Thomas, 1950.

Gallagher, Shaun and Morten Overgaard. "Introspection without Introspeculations," In *Pain: New Essays on Its Nature and the Methodology of Its Study*, edited by Murat Aydede, 277-290. Cambridge, MA: MIT Press, 2005.

Glennan, Stuart. "Rethinking Mechanistic Explanations," *Philosophy of Science*, 69 (2002): 342-353.

Glennan, Stuart, "Modeling Mechanisms," *Studies in History and Philosophy of Biological and Biomedical Sciences*, 36 (2005): 443-464.

Grahek, Nikola, "Objective and Subjective Aspects of Pain," *Philosophical Psychology*,

4 (1991): 249-267.

Grahek, Nikola. *Feeling Pain and Being in Pain*. 2nd ed. Cambridge, MA: MIT Press, 2007.

Griffiths, Paul E. *What Emotions Really Are*. London: The University of Chicago Press, 1997.

Gustafson, Don. "Categorizing Pain." In *Pain: New Essays on Its Nature and the Methodology of Its Study*, edited by Murat Aydede, 219-241. Cambridge, MA: MIT Press, 2005.

Hacking, Ian. "A Tradition of Natural Kinds," *Philosophical Studies,* 61 (1991): 109-126.

Hacking, Ian, "On Boyd," *Philosophical Studies*, 61 (1991): 149-154.

Hardcastle, Valerie Gray. *The Myth of Pain*. Cambridge, MA: MIT Press, 1999.

Hansson, Per, "Difficulties in stratifying neuropathic pain by mechanisms," *European Journal of Pain* 7 (2003): 353-357.

Head, Henry and Gordon Holmes, "Sensory Disturbances from Cerebral Lesions," Croonian Lectures at the Royal College of Physicians. June 13, 15, 20, and 27, 1911.

Hill, Anne, "Phantom Limb Pain: A Review of the Literature on Attributes and Potential Mechanisms," *Journal of Pain and Symptom Management* 17 (1999): 125-142.

Hill, Christopher S., "Ow! The Paradox of Pain," In *Pain: New Essays on Its Nature and the Methodology of Its Study*, edited by Murat Aydede, 75-98. Cambridge, MA: MIT Press, 2005.

Jensen, Troels S., and Ralf Baron, "Translation of symptoms and signs into mechanisms in neuropathic pain," *Pain* 102 (2003): 1-8.

Jensen, Troels Staehelin and Nanna Brix Finnerup, "Neuropathic pain: Peripheral and central mechanisms," *European Journal of Pain Supplements*, 3 (2009): 33-36.

Julius, David and Allan I. Basbaum, "Molecular mechanisms of nociception," *Nature* 413 (2001): 203-210.

Kahane, Guy. "Feeling Pain for the Very First Time: The Normative Knowledge Argument," *Philosophy and Phenomenological Research,* forthcoming.

Krane, Elliot J. and Lori B. Heller, "The Prevalence of Phantom Sensation and Pain in Pediatric Amputees," *Journal of Pain and Symptom Management*, 10 (1995): 21-29.

Kripke, Saul. *Naming and Necessity*. Boston: Harvard University Press. 1980.

Lariviere, William R., Sonya G. Wilson, Tinna M. Laughlin, Anna Kokayeff, Erin E. West, Seetal M. Adhikari, You Wan, and Jeffrey S. Mogil, "Heritability of nociception. III. Genetic relationships among commonly used assays of nociception and hypersensitivity," *Pain* 97 (2002): 75-86.

Le Bars, Daniel, "The whole body receptive field of dorsal horn multireceptive neurons," *Brain Research Reviews*, 40 (2002): 29-44.

Legrain, Valéry, Gian Domenico Iannetti, Léon Plaghki, and André Mouraux, "The pain matrix reloaded. A salience detection system for the body," *Progress in Neurobiology* 93 (2011): 111-124.

Lewis, David, "Psychophysical and Theoretical Identifications," edited by David M. Rosenthal, 204-210. Oxford University Press, 1991.

Lippe, Philipp M., Charles Brock, Jose David, Ronald Crossno and Stuart Gitlow, "The First National Pain Medicine Summit—Final Summary Report," *Pain Medicine* 11 (2010): 1447-1468.

Livingston, W.K., *Pain Mechanisms*. New York: Macmillan. 1943.

Loeser, J.D., "Low Back Pain," In *Bonica's Management of Pain*, 3rd edition, edited by J.D. Loeser, 1508-1565. Philadelphia: Lippincott Williams & Wilkins, 1994.

Luger, Nancy M., David B. Mach, Molly A. Sevcik and Patrick W. Mantyh, "Bone Cancer Pain: From Model to Mechanism to Therapy," *Journal of Pain and Symptom Management*, 29 (2005): S32-S46.

Machamer, Peter, Lindley Darden, and Carl F. Craver, "Thinking About Mechanisms," *Philosophy of Science, 67* (2000): 1-25.

Machelska, Halina, Julia K. Schopohl, Shaaban A. Mousa, Dominika Labuz, Michael Schäfer, and Christoph Stein, "Different mechanisms of intrinsic pain inhibition in early and late inflammation," *Journal of Neuroimmunology*, 141 (2003): 30-39.

Machery, Edouard. *Doing Without Concepts.* Oxford: Oxford University Press, 2009.

Mackie, J.L., "Locke's Anticipation of Kripke," *Analysis*, 34 (1974): 177-180.

MacPherson, Fiona, "Individuating the Senses," In *The Senses: Classical and*

Contemporary Readings, edited by Fiona MacPherson. Oxford: Oxford University Press, forthcoming.

Marina de Tommaso, Michele Sardaro, and Paolo Livrea, "Aesthetic Value of paintings affects pain thresholds," *Consciousness and Cognition,* 17 (2008): 1152-1162.

Marris, Emma, "More pain studies needed: Scientists struggle to fund work on animal pain and distress," *Nature*, 458 (2009): 394-395.

Maroux, André, Ana Diukova, Michael C. Lee, Richard G. Wise, and Gian Domenico Iannetti, "A multisensory investigation of the functional significance of the 'pain matrix'," *Neuro Image*, 54 (2011): 2237-2249.

McMahon, Stephen B., and Martin Koltzenburg eds. 2006. *Wall anThe d Melzack's Textbook of Pain*. 5th edition. Copyright © 2006 Churchill Livingstone, An Imprint of Elsevier.

Melzack, Ronald. *The Puzzle of Pain*. New York: Basic Books, Inc., 1973.

Melzack, Ronald and Patrick D. Wall. *The Challenge of Pain*. 2nd ed. 1988. Reprint, London: Penguin Books, 2008.

Melzack, Ronald, Patrick D. Wall and T.C. Tye, "Acute pain in an emergency clinic: latency of onset and descriptor patterns," *Pain,* 14 (1982): 33-43.

Melzack, Ronald, "Pain and the Neuromatrix in the Brain," *Journal of Dental Education* 65 (2001): 1378-1382.

Merskey, Harold, "Response to Editorial: New Perspectives on the Definition of Pain," *Pain* 67 (1996): 209.

Millikan, Ruth Garrett. "Historical Kinds and the 'Special Sciences,'" *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition,* 95 (1999), 45-65.

Mogil, JS, SG Wilson, K Bon, SE Lee, K Chung, P Raber, JO Pieper, HS Hain, JK Belknap, L Hubert, GI Elmer, JM Chung, and M Devor, "Heritability of nociception I: responses of 11 inbred mouse strains on 12 measures of nociception," *Pain* 88 (1999a): 67-82.

Mogil, JS, SG Wilson, K Bon, SE Lee, K Chung, P Raber, JO Pieper, HS Hain, JK Belknap, L Hubert, GI Elmer, JM Chung, and M Devor, "Heritability of nociception II. 'Types' of nociception revealed by genetic correlation analysis," *Pain* 80 (1999b): 83-93.

Mogil, Jeffrey S., Lei Yu, and Allan I. Basbaum, "Pain Genes?: Natural Variation and

Transgenic Mutants," *Annual Reviews of Neuroscience* 23 (2000): 777-811.

Mogil, Jeffrey S., "Animal models of pain: progress and challenges," *Nature Reviews Neuroscience,* 10 (2009): 283-294.

Murhpy, Dominic and Stephen Stich, "Griffiths, Elimination & Psychopathology," *Metascience*, forthcoming.

Nahmias, Eddy. "The Problem of Pain," In *Pain: New Essays on Its Nature and the Methodology of Its Study*, edited by Murat Aydede, 307-314. Cambridge, MA: MIT Press, 2005.

Neugebauer, Volker, Vasco Galhardo, Sabatino Maione and Sean C. Mackey, "Forebrain pain mechanisms, " *Brain Research Reviews*, 60 (2009): 226-242.

Osuntokun, B.O., E.L. Odeku, and L. Luzzato, "Congenital pain asymbolia and auditory imperception," *Journal of Neurology, Neurosurgery, and Psychiatry with Practical Neurology*, 31 (1968): 291-296.

Parsons, Terence, "Essentialism and Quantified Modal-Logic," in *Reference and Modality*, edited by Leonard Linsky, 73-87. Oxford University Press, 1971.

Pitcher, George, "The Awfulness of Pain," *The Journal of Philosophy,* 67 (1970): 481 492.

Price, Donald D and Murat Aydede, "The Experimental Use of Introspection in the Scientific Study of Pain and Its Integration with Third-Person Methodologies: The Experiential-Phenomenological Approach," In *Pain: New Essays on Its Nature and the Methodology of Its Study*, edited by Murat Aydede, 75-98. Cambridge, MA: MIT Press, 2005.

Price, Donald D., James J. Barrell and Pierre Rainville, "Integrating experiential phenomenological methods and neuroscience to study neural mechanisms of pain and consciousness," *Consciousness and Cognition*, 11(2002): 598-608.

Putnam, Hilary. "The meaning of `meaning"." In K. Gunderson (Ed.), Language, Mind and Knowledge, Vol. VII of Minnesota Studies in the Philosophy of Science, 131-193. University of Minnesota Press, 1975.

Putnam, Hilary, "The Nature of Mental States," In *The Nature of Mind*, edited by David M. Rosenthal, 197-203. Oxford University Press, 1991.

Quine, W.V., "Natural Kinds," In *Naming, Necessity, and Natural Kinds*, edited by Stephen P. Schwartz, 155-175. London: Cornell University Press, 1979.

Radden, Jennifer. *Moody Minds Distempered: Essays on Melancholy and Depression*. New York: Oxford University Press, 2009.

Rainville, P., G.H. Duncan, D.D. Price, B. Carrier, and M.C. Bushnell. "Pain Affect Encoded in Human Anterior Cingulate but not Somatosensory Cortex." *Science* 277 (1997): 968-971.

Ramachandran, V.S. and William Hirstein, "The Perception of Phantom Limbs: The D.O. Hebb Lecture," *Brain,* 121 (1998): 1603-1630

Rose, James D., "The Neurobehavioral Nature of Fishes and the Question of Awareness and Pain," *Reviews in Fisheries Science* 10 (2002): 1-38.

Rosenthal, David M. "Sensory Qualities, Consciousness, and Perception," In *Consciousness and Mind,* 175-226. *Oxford: Clarendon Press.* (2005).

Rubins, Jack L., and Emanuel D. Friedman, "Asymbolia for Pain," *Archives of Neurology and. Psychiatry,* 60 (1948): 554-573.

Ruse, Michael, "Darwinism and mechanism: metaphor in science," *Studies in History and Philosophy of Biological and Biomedical Sciences*, 36 (2005): 285-302.

Scholz, Joachim and Clifford J. Woolf, "Can we conquer pain?," *Nature Publishing Group*, 5 (2002): 1062-1067.

Schwitzgebel, Eric, "The Unreliability of Naive Introspection," *Philosophical Review*, 117 (2008): 245-273.

Shriver, Adam, "Knocking out Pain in Livestock: Can Technology Succeed Where Morality has Stalled?," *Neuroethics,* 2 (2009): 115-124.

Skipper Jr., Robert A., and Robert L. Millstein, "Thinking about evolutionary mechanisms: natural selection," *Studies in History and Philosophy of Biological and Biomedical Sciences*, 36 (2005): 327-347.

Smart, Keith and Catherine Doody, "Mechanism-based clinical reasoning of pain by experienced musculoskeletal physiotherapists," *Physiotherapy* (2006): 171-178.

Smart, Keith M., Neil E. O'Connell and Catherine Doody, "Towards a mechanisms based classification of pain in musculoskeletal physiotherapy?," *Physical Therapy Reviews*, 13, no. 1(2008): 1-10.

Staud, Roland, Charles J. Vierck, Richard L. Cannon, Andre P. Mauderli, and Donald D. Price, "Abnormal sensitization and temporal summation of second pain (wind-up) in patients with fibromyalgia syndrome," *Pain* 92 (2001): 165-

175.

Thorn, Beverly E. and John W. Burns, "Common and specific treatment mechanisms in psychosocial pain interventions: the need for a new research agenda," *Pain*, forthcoming, 2011.

Tommerdahl, M., K.A. Delemos, C.J. Vierck, Jr., O.V. Favorov, and B.L. Whitsel, "Anterior Parietal Cortical Response to Tactile and Skin-Heating Stimuli Applied to the Same Skin Sit," *Journal of Neurophysiology*, 75 (1996): 2662 2670.

Tommerdahl, M., K.A. Delemos, O.V. Favorov, C.B. Metz, C.J. Vierck, Jr., and B.L. Whitsel, "Response of Anterior Parietal Cortex to Different Modes of Same-Site Skin Stimulation," *Journal of Neurophysiology* 80 (1998): 3272-3283.

Tye, Michael. "Another Look at Representationalism about Pain," In *Pain: New Essays on Its Nature and the Methodology of Its Study*, edited by Murat Aydede, 99-120. Cambridge, MA: MIT Press, 2005.

Vilayanur S. Ramachandran, Paul D. McGeoch, Lisa Williams and Gerard Arcilla, "Rapid Relief of Thalamic Pain Syndrome Induced by Vestibular Caloric Stimulation," *Neurocase*, 13 (2007): 185-188.

Villemure, Chantal, Burton M. Slotnick and M. Catherine Bushnell, "Effects of odors on pain perception: deciphering the roles of emotion and attention," *Pain* 106 (2003): 101-108.

Wall, P.D., "Response to Editoral by Anand and Craig," *Pain* 67 (1996): 209.

Wade, James B., Linda M. Dougherty, C. Ray Archer, and Donald D. Price, "Assessing the stages of pain processing: a multivariate analytical approach," *Pain* 68 (1996): 157-167.

Weiskopf, Daniel Aaron, "The plurality of concepts," *Synthese* 169 (2009): 145-173.

Westlake, E.K., "Congenital Indifference to Pain," *British Medical Journal*, (1952): 144.

Woolf, Clifford J., Gary J. Bennett, Michael Doherty, Ronald Dubner, Bruce Kidd, Martin Koltzenburg, Richard Lipton, John D. Loeser, Richard Payne, and Eric Torebjork, "Towards a mechanism-based classification of pain?," *Pain*, 77 (1998): 227-229.

Woolf, Clifford J. and Isabelle Decosterd, "Implications of recent advances in understanding of pain pathophysiology for the assessment of pain in patients," *Pain* 6 (1999): S141-S147.

Woolf, Clifford J. and Richard J. Mannion, “Neuropathic pain: aetiology, symptoms, mechanisms, and management,” *The Lancet* (1999): 1959-1964.

Woolf, Clifford J., and Michael W. Salter, “Neuronal Plasticity: Increasing the Gain in Pain,” *Science* 288 (2000): 1765-1768.

Woolf, Clifford J. and Mitchell B. Max, “Mechanism-based Pain Diagnosis,” *Anesthesiology* (2001): 241-249.

Zhang, Xuming and Peter A. McNaughton. "Why Pain Gets Worse: The Mechanism of Heat Hyperalgesia," *Journal of General Physiology*, 128 no.5 (2006): 491- 493.

Ziv-Sefer, Sagit, Pnina Raber, Shahar Barbash, and Marshall Devor, “Unity vs. diversity of neuropathic pain mechanisms: Allodynia and hyperalgesia in rats selected for heritable predisposition to spontaneous pain,” *Pain* 146 (2009): 148-157.