

PAIN LINGUISTICS: A CASE FOR PLURALISM

BY SABRINA CONINX,¹ PASCALE WILLEMSSEN² AND KEVIN REUTER²

The most common approach to understanding the semantics of the concept of pain is third-person thought experiments. By contrast, the most frequent and most relevant uses of the folk concept of pain are from a first-person perspective in conversational settings. In this paper, we use a set of linguistic tools to systematically explore the semantics of what people communicate when reporting pain from a first-person perspective. Our results suggest that only a pluralistic view can do justice to the way we talk about pain from a first-person perspective: The semantic content of the folk concept of pain consists of information about both an unpleasant feeling and a disruptive bodily state. Pain linguistics thus provides new insights into ordinary pain language and poses an interesting challenge to the dominant unitary views of pain.

Keywords: folk concept of pain, bodily states, feeling pain, paradox of pain, deniability test, projection test.

I. INTRODUCTION

What is pain? In the philosophical tradition, this question has received a rather univocal answer. For example, according to Lewis (1980: 222), ‘Pain is a feeling. Surely that is uncontroversial. To have pain and to feel pain are one and the same.’ This understanding of pain also appears to prevail in the medical sciences. Prominently, the International Association for the Study of Pain defines pain as ‘an unpleasant sensory and emotional experience’ (Raja *et al.* 2020: 1977). These definitions do not try to introduce a technical term for the sole purpose of philosophical or scientific discussion but aim to capture the *folk concept of pain* (Goldberg, Reuter and Sytsma 2023). However, the folk concept of pain is far less uniform than most common definitions suggest.

Hill (2005) argues that the folk concept of pain appears to be oddly paradoxical, as it pulls into two directions that seem to mutually exclude each other. On the one hand, people often treat pain as a subjective and private feeling, commonly characterized as being unpleasant or hurtful. On the other hand,

people often treat pain as an objective and publicly accessible bodily state, widely identified as physical damage, disruption, or disturbance. This apparent paradox not only complicates the identification of the research subject for philosophical and scientific investigations (Coninx 2020), but also indicates potential ambiguities in everyday language (Liu 2021; Sytsma and Reuter 2017; Salomons *et al.* 2021). These ambiguities might constitute a severe source of miscommunication among medical staff, health care providers, caregivers, and laypeople when reporting pain. Thus, understanding the folk concept of pain and whether or not it is paradoxical has significant implications for understanding, treating, and interacting with pain patients.

In this paper, we empirically investigate the folk concept of pain by examining how people understand *first-person pain reports*, such as ‘I have a pain in my arm.’ More specifically, our objective is to identify content that is semantically encoded as part of the concept of pain and content that is only conversationally implicated. To this end, we adapted three linguistic tests—the implication, projection, and deniability test—and applied them in two preregistered experiments. The results of our experiments align with the predictions of *pluralistic views*. That is, the best explanation of our findings is that paradigmatic first-person pain reports semantically entail information about both an unpleasant feeling and a disruptive bodily state. These results stand in contrast to unitary views that predominate the recent literature and only consider one of these aspects as part of the semantic content of the folk concept of pain.

The paper is structured as follows. In Section II, we introduce the most prominent views on the folk concept of pain and critically discuss recent methodological approaches. Section III introduces our new experimental-linguistic approach, while Section IV presents the results of our first preregistered experiment using the implication and projection tests. Section V provides the results of the second experiment using the deniability test. Based on the results of these three tests, we argue for a complex, pluralistic view in Section VI. Information regarding both a bodily disruption and an unpleasant feeling is likely to be communicated as part of the semantic content of the folk concept of pain. Furthermore, we discuss the potential methodological limitations of our design and the plausibility of different interpretations of the pluralistic view. Section VII concludes our results.

II. RECENT APPROACHES TO THE FOLK CONCEPT OF PAIN

Unitary views have been at the centre of the debate concerning the folk concept of pain for the past decade. Unitary views assume that there is a single folk concept of pain and that it has a distinct univocal meaning, referring either to a feeling or a bodily state.

According to the *feeling view*, people commonly treat pain as a private and subjective feeling (Aydede 2009; Tye 2017). Some authors have characterized

this aspect as referring to the subject's mental state more generally (Borg *et al.* 2020, 2021; Liu 2022). However, the typical way of thinking about pain as a mental state is to treat it as a conscious feeling (Borg *et al.* 2021; Kripke 1981; Lewis 1980; Tye 2017) that is unpleasant or hurtful, at least in paradigmatic cases (Coninx 2022; Raja *et al.* 2020).

By contrast, according to the *bodily view*, people commonly treat pain as a public and objective bodily state (Kim *et al.* 2016; Massin 2017; Reuter, Phillips and Sytsma 2014; Reuter and Sytsma 2020; Sienhold and Sytsma 2019). Some authors have assumed that this aspect refers to the state of a body part without further specifications (Hyman 2003; Liu 2022). However, the typical way of thinking about pain as a bodily state is to treat it as a particular physical condition of the (non-brain based) body (Reuter 2017; Salomons *et al.* 2021). While there is little agreement regarding the precise definition of this condition, most suggestions revolve around the idea of there being something physically wrong with a body part, as in the case of bodily damage, disturbance, or disruption (Borg *et al.* 2020; Reuter and Sytsma 2020; Salomons *et al.* 2021).

Both unitary views have been criticized in the more recent literature for telling only one part of the story: People appear to be willing to treat pain as a feeling *and* as a bodily state (Borg *et al.* 2020; Reuter and Sytsma 2020; Liu 2021, 2022; Salomons *et al.* 2021). These considerations gave rise to different versions of the *pluralist view* (Borg, Hansen and Salomons 2019; Borg *et al.* 2020; Liu 2021, 2022). Pluralist views are united in rejecting the assumption of a (single) folk concept of pain with a univocal meaning: They all claim that the folk concept of pain is not always treated as referring only to an unpleasant feeling or only to a disruptive bodily state. Instead, the concept is assumed to be more complex than the different versions of the unitary view would have us believe (for a more detailed discussion of different versions of the pluralist view, see Section VI.3).

In 2010, experimental philosophers began to investigate the folk concept of pain empirically via *vignette studies*.¹ These studies mainly serve the purpose of deciding between the two unitary views, namely the feeling view and the bodily view. As a result, they typically describe scenarios in which the agent either experiences a feeling or undergoes some bodily disturbance, whereas the other aspect is absent (Reuter, Phillips and Sytsma 2014; Reuter and Sytsma 2020; Salomons *et al.* 2021). However, the empirical evidence is mixed, with some studies supporting the feeling view and others the bodily view.

Several studies have challenged the assumption that the folk concept of pain refers to a subjective experience (Reuter 2011; Reuter, Phillips and Sytsma 2014;

¹ While vignette-based and corpus-based studies have dominated the empirical investigations into the folk concept of pain, Liu (2021) used linguistic methods to test for ambiguity in pain-related words such as 'sore', 'aching', and 'hurting'. However, these tests cannot be performed directly for the term 'pain'.

Reuter and Sytsma 2020; Sytsma 2010; Sytsma and Machery 2009). One direct consequence of this assumption is that pains cannot exist unfelt. However, participants in several experimental studies were willing to ascribe pain to people in the absence of a corresponding feeling. For example, a majority of people believe a severely wounded soldier has pain even if he does not *feel* any pain (Reuter and Sytsma 2020). This finding might motivate the conclusion that the folk concept of pain only entails bodily aspects.

Other authors have challenged the assumption that the folk concept of pain (always) refers to a bodily state (Borg *et al.* 2020; Salomons *et al.* 2021). According to this assumption, pains cannot exist when there is no physical damage, disturbance, or disruption. However, in experimental studies, participants were willing to ascribe pain to people lacking a corresponding bodily state. For example, a majority of people ascribed pain to people who reported feeling pain as a result of the direct stimulation of their brain without any (non-brain based) bodily changes taking place (Salomons *et al.* 2021).

These results indicate that, when participants were presented with different scenarios, they were willing to ascribe pain even in the absence of an unpleasant feeling or a disruptive body state. Depending on the context provided and the amount of detail in which it is described, people treat pain as a feeling in some cases and as a bodily state in others. A possible explanation for these results is that the folk concept of pain is inherently messy and can be ‘pushed’ around quite easily.

There is another way to think about the variability in people’s judgements. We believe that well-documented experimental-pragmatic effects could partly explain the contrasting results. It has been argued that ‘pulling apart features that usually go together’, a practice quite common in experimental-philosophical vignette studies, can cause confusion in participants and affect their conceptual competence (Machery 2017: 117). While people are perfectly competent to apply the term ‘pain’ in ordinary cases in which feeling and bodily aspects co-occur, they rarely experience cases in which one of the two is absent. However, most vignette studies on the folk concept of pain do exactly that. Furthermore, Machery (2017) identified several other characteristics that might undermine the reliability of judgements that would otherwise be reliable. For example, while people are familiar with bruised knees, cut fingers, or headaches, few have been confronted with soldiers who have gunshot wounds but do not feel any pain. Thus, while many of the vignettes used in that research were aimed at pushing a concept to its limits, they may have pushed too hard. Finally, vignette-based research on pain introduces an additional level of unfamiliarity. Participants are asked to ascribe pain to an unknown person from a third-person perspective without being involved or engaging with said person: The attribution of pain is thus taken out of the social context in which it normally takes place and might have meaningful consequences. In everyday life, we rarely consider and communicate to others whether we are

willing to attribute pain to an (absent) third person in this manner. Instead, in real-world cases, we seem more often confronted with contexts in which we ascribe pain to ourselves or to people in direct interaction (e.g. communication with children, partner, and patients).

To reduce the influence of these potential sources of distraction, we suggest a new approach to the folk concept of pain by focusing on pain *linguistics*. Instead of creating vignettes to address our concerns, we decided to have a speaker utter a first-person pain report and ask participants what they inferred from such a statement. Situations in which someone says, ‘I have a pain in my arm,’ are assumed to be more familiar to the participants and thus reliably trigger attempts to make sense of this statement, similar to social interaction in everyday life. Crucially, our participants were not asked to judge whether the person had a pain. The person says that they do. What is relevant for us is what our participants infer from such a statement.

We consider the linguistic studies of first-person pain reports to be an innovative addition to our empirical toolbox. However, we do not suggest that pain linguists should replace vignette-based and corpus-based research. In Section VI, we discuss the compatibility of our results with vignette studies, and we present methodological and theoretical implications.

III. A NEW APPROACH: IMPLICATION, PROJECTION, AND DENIABILITY

All attempts to understand the folk concept of pain are united by the search for the *semantic features* of this folk concept. A promising approach, therefore, begins with identifying the features reliably conveyed when a speaker uses the term ‘pain’, and then differentiating the semantic features from those that are merely conveyed pragmatically. On all accounts—feeling views, bodily views, or pluralist views—first-person pain reports such as, ‘I have a pain in my arm,’ reliably convey two pieces of information (in the remainder of this paper, we refer to such implied content as *implications*):

[Feeling] The speaker has an unpleasant feeling.

[Bodily] Something is physically wrong with the speaker’s body (or the speaker at least believes that something is physically wrong with their body).

Under normal circumstances, we consider first-person pain reports to express both types of information. However, the fact that these features are conveyed in seemingly all ordinary uses does not provide sufficient evidence that they are *semantically entailed*. Implications can be communicated in three different ways by a target statement: They can be (1) semantically entailed,

(2) presupposed, or (3) conversationally implicated.² Consider the often-used regret case to illustrate the three different aspects:

[Target Statement] 'I regret drinking instant coffee this morning.'³

Let us further consider the following potential *implications*:

- (a) Tom has a negative feeling about having drunk instant coffee.
- (b) Tom drank instant coffee this morning.
- (c) Tom prefers freshly brewed coffee.
- (d) Tom likes flowers.

Obviously, (d) cannot be inferred from the target statement and is, therefore, not an implication. However, (a) to (c) can be inferred.⁴ Claim (a) partly expresses what is literally meant by Tom's statement; that is, (a) is *semantically entailed* by the target statement. Claim (b) does not express what is stated literally by Tom's statement but rather what is presupposed by it. If Tom had not drunk instant coffee, then it would not even make sense to say that he regretted drinking coffee. Finally, claim (c) neither expresses what is literally meant by Tom's statement nor what is presupposed; it is *conversationally implicated*. Depending on the context, this inference can be made, but the inference is beyond what has literally been said. Linguists have devised several tests providing evidence for whether a certain content is semantically entailed, presupposed, or conversationally implicated.

Implication Test. Before we can discuss how a piece of information is conveyed, we need to determine whether it has actually been conveyed at all. For example, the information, 'Tom likes flowers,' cannot be inferred from the target statement. Therefore, it is not an implication of the utterance, 'I regret drinking instant coffee this morning.' The implication test determines what pieces of information are implications of a target statement. For this purpose, people answer a simple question such as, 'From this statement alone and having no other information, what do you infer from this statement?' Semantically

² There is a fourth way in which implications can be conveyed, namely by means of conventional implicatures. Conventional implicatures are typically considered part of the conventional meaning of the words that convey them while not being part of their truth-conditional meaning. We do not consider it plausible and are unaware of any suggestions that pain conventionally implicates bodily and/or feeling content. It should also be mentioned that some scholars have been explicitly sceptical of the entire class of conventional implicatures and its usefulness in philosophical and linguistic theorizing (e.g. Bach 2006; Zakkou 2018).

³ We have borrowed this example verbatim from Väyrynen (2013: 60). Similar examples can also be found in the philosophical and linguistic literature in the study of especially presuppositions (e.g. Horn 1986).

⁴ Some scholars have denied that semantically entailed content is ever *inferred* and claimed that it is given to us directly (Austin 1962; Grice 1975). As semantically entailed content is conveyed at the level of what-is-said explicitly, there is no need for an additional inferential process. Our use of the term 'inferred' is less technical and more colloquial. It is intended to mean 'what one understands upon hearing the utterance.'

entailed, as well as presupposed content, should always be inferred by competent speakers who understand the meaning of the statement and the terms involved. Conversational implicatures are also often inferred, but the extent to which they are inferred depends on the strength of the conversational implicature. While particularized conversational implicatures depend strongly on context and, thus, are often less strongly inferred, generalized conversational implicatures are more strongly inferred due to their independence from the context (for a more detailed discussion and illustrations, see e.g. Davis 2019).

Projection Test. The projection test helps to determine whether an implication is semantically entailed or is presupposed (Chierchia and McConnell-Ginet 2000; Huang 2006; Levinson 1983) by embedding the target statement in an entailment-cancelling operator, such as a negation. Here is an example:

[Target Statement] ‘I regret drinking instant coffee this morning.’

[Negation] ‘I do not regret drinking instant coffee this morning.’

The entailment-cancelling operator does precisely what the name suggests: It cancels semantically entailed content. If Tom *negates* that he regrets drinking instant coffee, then he certainly no longer conveys the information that he has a negative feeling about it. In contrast, however, the presupposed content, namely that he drank coffee in the morning, survives. Regardless of whether or not Tom regrets drinking instant coffee, we can still infer that he must have drunk coffee.

Deniability Test. Finally, the deniability test provides a valuable tool to distinguish semantically entailed and presupposed content from conversationally implicated content. In the deniability test, participants are asked how contradictory it sounds to explicitly deny or cancel certain content (Reins and Wiegmann 2021). Other than semantically entailed or presupposed content, conversational implicatures can (often) be denied (Blome-Tillmann 2008; Zakkou 2018). If Tom were to say, ‘I regret drinking instant coffee this morning,’ and then denied that he preferred freshly brewed coffee, then he would not contradict himself. By contrast, if Tom said, ‘I regret drinking instant coffee this morning,’ but denied that he had a negative feeling about drinking instant coffee, then his statement would be indeed contradictory.

Note that the deniability test is closely related to the more common cancellability test (Davis 2019; Grice 1975; Zakkou 2018). The cancellability test asks whether an original statement that triggers an implication can be combined with the immediate and explicit denial of that implication; for example, ‘This is round, but I do not mean to say that it has no edges.’ This test has been applied successfully in different experimental studies (Almeida, Struchiner and Hannikainen 2023; Baumgartner, Willemsen and Reuter 2022; Muth, Briesen and Carbon 2020; Sytsma, Willemsen and Reuter 2022; Willemsen and Reuter 2021). The deniability test also investigates whether the speaker can take back an implication, but it does so in more elaborate

Table 1. Prediction for implication, projection, and deniability for semantic entailment (SE), presupposition (Presup), and conversational implicature (CI).

	Implication	Projection	Deniability
Semantic entailment	☑	☒	☒
Presupposition	☑	☑	☒
Conversational implicature	☑	☑/☒	☑

conversational settings.⁵ For the purpose of experimental studies, the deniability test usually provides a more natural conversational context, such as in communication between patients and medical professionals, which we have indicated as being central to the understanding of the folk concept of pain.

As indicated in Table 1, combining these three traditional linguistic tests provides a checklist that might help identifying semantically entailed content, presupposed content, and conversationally implicated content. On that basis, we will categorize a piece of information as being semantically entailed if the following conditions are fulfilled: (i) The content is inferred from the target statement (implication test), (ii) the content does not project under an entailment-cancelling operator (projection test), and (iii) the content is not deniable (deniability test).

Based on this, we can return to the first-person pain report, ‘I have a pain in my arm,’ as our target statement. The implication, projection, and deniability tests should allow us to determine whether empirical data better support the feeling, bodily, or pluralist views. To do so, we need to define different contents that adequately represent the respective feeling and bodily aspects. We suggest the following three target contents:

Body₁: There is something physically wrong with Tom’s arm.

Body₂: Tom thinks that there is something physically wrong with his arm.

Feeling: Tom feels something unpleasant.

The exact phrasing of all three target contents is, of course, debatable. We decided to employ a commonly accepted characterization of the mental aspect as an unpleasant feeling. Representing the bodily view accurately was yet more challenging, as to be discussed in the following.

At the core of bodily views seems to be the general description that there is *something physically wrong* with a person’s body (see Section II). However, one might think that this interpretation comes in two versions: a more objective one (*Body₁*) and a more subjective one (*Body₂*). *Body₁* represents an objective

⁵ The deniability test has recently been applied by Reins & Wiegmann (2021) in investigating the folk concept of lying. When examining the plausibility of contextualist and relativist intuitions concerning taste predicates, Kneer (2021, 2022) uses a comparable experimental design, which he calls *retractability test*.

interpretation, stating that there is, as a matter of fact, something physically wrong with Tom's arm. However, it might be objected that this matter of fact is not essential. What is essential to the bodily view is that the speaker subjectively *thinks* or *believes* that something is wrong with his arm, at least when it comes to the semantic content of first-person pain report, such as 'I have a pain in my arm.' This understanding is expressed by *Body₂*. *Body₁* and *Body₂* are alternative readings, and we do not assume that defenders of the bodily view need to necessarily accept both, while they need to accept either an objective or a subjective interpretation of the bodily view.

Independently of whether defenders of the bodily view accept a subjective or objective interpretation, one might still argue that the concrete wording of *Body₁* and *Body₂* is questionable, as our formulations draw attention to one specific body part, namely the speaker's *arm*. However, the bodily view is in principle compatible with a more inclusive understanding concerning a person's body as such—that might also account for *referred pains* in which the location of a bodily disturbance and the reported location dissociate. These special cases and related limitations of our methodological design are discussed in more detail in Section VI.2.

We are now in a position to state the predictions that the bodily, feeling, and pluralist views would make with regard to the first-person pain statement 'I have a pain in my arm.'

Bodily View: Whereas *Body₁* (and *Body₂*) is semantically entailed, *Feeling* is (at best) conversationally implicated by the target statement 'I have a pain in my arm.'

Feeling View: Whereas *Feeling* is semantically entailed, *Body₁* (and *Body₂*) is (at best) conversationally implicated by the target statement 'I have a pain in my arm.'

Pluralist View: Both *Body₁* (and *Body₂*) and *Feeling* are semantically entailed by the target statement 'I have a pain in my arm.'

As the implication and projection tests are closely related, we decided to run them together in one experiment, the results of which we present in Section IV. The results of the deniability test are presented separately in Section V. To ensure that the experiments were well-designed, we included the regret condition as a control condition.

IV. STUDY 1: IMPLICATION AND PROJECTION

Our new experimental approach applies traditional linguistic tests to study the semantics of first-person pain reports. Instead of presenting participants with vignette-based stimuli that are prone to various contextual biases, we presented the participants with a single-sentence stimulus, 'I have a pain

in my arm.’ This enabled us to focus on natural conversational contexts in which people ascribed pain to themselves and reported this judgement to others instead of asking participants to determine whether an agent had pain from a third-person perspective. The aim of Study 1 was two-fold. First, we examined whether various potential implications that might be triggered by a pain statement were inferred in a positive embedding. Secondly, we tested which of these implications were retained in a negative embedding; that is, under negation.

IV.1 Methods

We implemented a $2 \times 2 \times 6$ mixed design with the between-subject factors *embedding* (positive, negative) and *concept* (pain, regret), and the within-subject factor *implication*. This means that we manipulated (a) whether the formulation of the first-person report was positive or contained a negation and (b) whether that first-person report was about the person having pain in their arm or regretting drinking instant coffee. This leads to four conditions, and participants only received one of those four conditions. We asked the participants to imagine our protagonist Tom making one of the following four first person reports:

Pain Positive: ‘I have a pain in my arm.’

Pain Negative: ‘I don’t have a pain in my arm.’

Regret Positive: ‘I regret drinking instant coffee this morning.’

Regret Negative: ‘I don’t regret drinking instant coffee this morning.’

The third, *within-subject factor* consisted of six implications in the case of pain and five implications in the case of regret, following the first-person report. After the prompt, ‘From this statement alone and having no other information, what do you infer from this statement?’ participants received the following six implications for the pain conditions in randomized order:

Body1: There is something physically wrong with Tom’s arm.

Body2: Tom thinks that there is something physically wrong with his arm.

Feeling: Tom feels something unpleasant.

CI Pain: Tom needs help.

Presup Pain: Tom has an arm.

Unrelated: Tom likes flowers.

In the regret condition, subjects were presented with the following five statements in randomized order:

Neg Feeling: Tom has a negative feeling about drinking instant coffee this morning.

Wish: Tom wishes he had not drunk instant coffee this morning.

CI_Regret: Tom prefers freshly brewed coffee.

Presup_Regret: Tom drank instant coffee this morning.

Unrelated: Tom likes flowers.

The positive embeddings (*Pain Positive* and *Regret Positive*) were used in the implication test to determine the content that could be inferred based on the original pain or regret statement. The negative embeddings (*Pain Negative* and *Regret Negative*) were used in the projection test to investigate projection behaviour and to thus differentiate content that was presupposed from that which was not.

Body₁, *Body₂*, and *Feeling* were the target contents that were used to test the plausibility of the bodily view, the feeling view, and the pluralist view regarding the folk concept of pain. *Neg_Feeling* and *Wish* functioned as contents that were likely to be semantically entailed by the regret statement, allowing us to evaluate whether our experiments were designed in such a way that enabled us to identify the semantic features of a target statement. *Unrelated* functions in the pain and regret conditions served as a control for the implication test to exclude content that was not inferred. *Presup_Pain* and *Presup_Regret* were likely to be presupposed by the target statement. In contrast to the semantically entailed content, they should survive in the projection test. *CI_Pain* and *CI_Regret* are inferences that could be made based on the respective target statements, but which are inferred beyond what is literally said. Accordingly, it should be possible to deny them without producing a contradiction. These conversational implicatures will be the focus of the deniability test in Section V.

IV.2 Preregistered hypotheses

We preregistered (<https://osf.io/6jvs5>) the following hypotheses for the implication and projection tests:

- H₁:** For the contents *Body₁*, *Body₂*, and *Feeling*, as well as *Neg_Feeling* and *Wish*, the ratings are significantly above the midpoint of 5 for the positive embeddings and significantly below the midpoint for the negative embeddings.
- H₂:** For the contents *Presup_Pain* and *Presup_Regret*, the ratings are significantly above the midpoint of 5 for both positive and negative embeddings.
- H₃:** For *Unrelated*, the ratings are significantly below the midpoint of 5 for both positive and negative embeddings in the pain and regret condition.

A total of 262 participants were recruited via Prolific. All participants were at least 18 years old, native English speakers (or bilingual), and had an approval rate of at least 95 per cent. The participants had an average age of 38.47 years, and the gender distribution in the sample was 115 males, 141 females, and six non-binary persons.

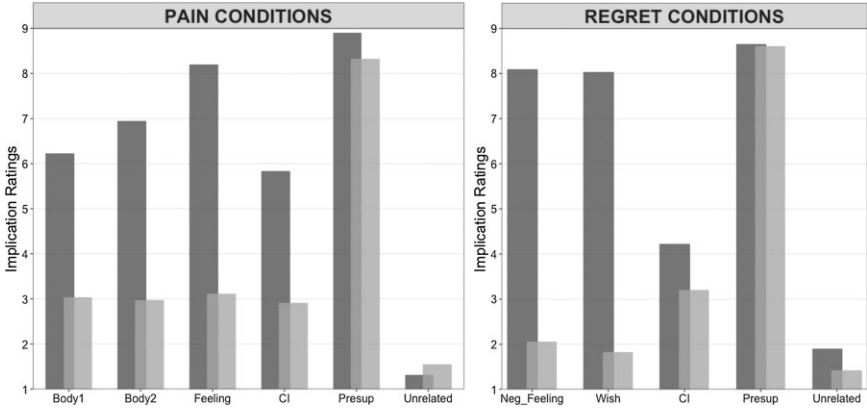


Figure 1. Mean implication ratings for the six pain statements (left) and five regret statements (right). The dark grey bars indicate the ratings for the positive conditions (implication test), and the light grey bars for the negative conditions (projection test).

IV.3 Results

The results of the implication and projection tests are depicted for the pain conditions in Fig. 1 (left) and Table 2 (Pain Positive and Pain Negative) and for the regret conditions in Fig. 1 (right) and Table 3 (Regret Positive and Regret Negative). The figures and tables depict the mean implications ratings for all implications in the pain and regret conditions, and for both the implication

Table 2. Mean implication ratings for the six pain statements in the positive conditions (implication test, upper part) and the negative conditions (projection test, lower part) and results of *t*-tests against the midpoint of the scale.

	Condition	Mean	Standard error	<i>t</i>	<i>p</i> -value
Implication (positive)	Body1	6.22	0.291	4.210	<0.001
	Body2	6.94	0.258	7.528	<0.001
	Feeling	8.19	0.149	21.403	<0.001
	CI.Pain	5.84	0.260	3.213	=0.002
	Presup.Pain	8.90	0.043	90.086	<0.001
	Unrelated	1.31	0.161	-22.918	<0.001
Projection (negative)	Body1	3.03	0.304	-6.486	<0.001
	Body2	2.97	0.301	-6.733	<0.001
	Feeling	3.12	0.331	-5.690	<0.001
	CI.Pain	2.91	0.307	-6.801	<0.001
	Presup.Pain	8.33	0.186	17.907	<0.001
	Unrelated	1.55	0.218	-15.800	<0.001

Table 3. Mean implication ratings for the five regret statements in the positive conditions (implication test, upper part) and the negative conditions (projection test, lower part) and results of *t*-tests against the midpoint of the scale.

	Condition	Mean	Standard error	<i>t</i>	<i>p</i> -value
Implication (positive)	Neg_Feeling	8.10	0.213	14.588	<0.001
	Wish	8.04	0.227	13.443	<0.001
	CI_Regret	4.22	0.318	-2.439	=0.017
	Presup_Regret	8.66	0.145	25.201	<0.001
	Unrelated	1.90	0.272	-11.429	<0.001
Projection (negative)	Neg_Feeling	2.06	0.233	-12.592	<0.001
	Wish	1.83	0.221	-14.341	<0.001
	CI_Regret	3.20	0.302	-5.961	<0.001
	Presup_Regret	8.61	0.156	23.187	<0.001
	Unrelated	1.42	0.186	-19.244	<0.001

and the projection test. In addition, the tables present the results of one-sample *t*-tests against the midpoint of 5. We conducted these tests to test our hypotheses **H1**, **H2**, and **H3**. Our results confirmed **H1**: All five statements (*Body1*, *Body2*, *Feeling*, *Neg_Feeling*, and *Wish*) received ratings that were significantly above the midpoint for the positive embedding and below the midpoint for the negative embedding. The two presuppositions (*Presup_Pain* and *Presup_Regret*) received ratings above the midpoint for the positive and negative claims, thus providing strong evidence for **H2**. The ratings for the unrelated statement were below the midpoint for both embeddings in the pain and regret conditions, thus supporting **H3**.

Post-hoc pairwise comparisons (not preregistered) between *Body1* and *Body2* did not reveal a significant difference ($p = 0.309$), while both bodily conditions were significantly different from the *Feeling* condition (both $ps < 0.001$).⁶

IV.4 Discussion

In the pain condition, we investigated three candidates (*Body1*, *Body2*, and *Feeling*) as being potentially communicated by means of semantic entailment. The data suggest that bodily (*Body1* and *Body2*) and feeling aspects (*Feeling*) were implicated but not presupposed by the claim ‘I have a pain in my arm’ because they did not project under negation. Furthermore, the putatively presupposed content (*Presup_Pain*) was rated as expected: In the positive and negative embed-

⁶ A repeated-measures ANOVA to examine the effects in the positive pain condition yielded a statistically significant result for the within-subjects factor *implication* ($F(1,65) = 102.24, p < 0.001, \eta^2 = 0.61$).

dings, the participants assumed that Tom had an arm in order to make sense of the target statement. Minor differences in the ratings for the positive and negative embeddings could be explained by experimental-pragmatic factors. Unrelated content was not implicated by the pain statement. Although we did not make hypotheses concerning the putatively conversationally implicated content (*CI_Pain*), it should be noted that it received ratings significantly above the midpoint for the positive condition and ratings significantly below the midpoint for the negative condition. The regret condition functioned as expected. Ratings for the semantically entailed contents in the regret condition (*Neg_Feeling* and *Wish*) were higher in the positive embedding and lower in the negative embedding compared with the potentially semantically entailed contents in the pain condition (*Feeling*, *Body1*, and *Body2*). This is not surprising. The regret case is a paradigmatic example indicating the experimental design's general functionality. Given the existing debate surrounding the folk concept of pain (see Section II), more ambivalent results are to be expected in the pain condition.

V. STUDY 2: DENIABILITY

In Study 1, we found that body and feeling features were implications of first-person pain reports that were not presupposed. Together with the conversational implicature contents, these conditions were transferred to Study 2. We did not include the presupposed contents (*Presup_Pain* and *Presup_Regret*) and unrelated content (*Unrelated*) from Study 1 because the previously presented results had determined their identities. Study 2 uses a variation of the cancellability test, which is also known as the deniability test. We adapted the paradigm by creating a new version of Reins & Wiegmann's (2021) deniability test. The deniability paradigm is particularly useful because it is discursive. This is a natural setting for investigating pain statements, as it is the type of communication situation that patients and doctors typically encounter.

V.1 Methods

We implemented a 7×1 between-subjects design with implication as a between-subjects factor. This means that participants were randomly assigned to one of seven conversations related to *Body1*, *Body2*, *Feeling*, *CI_Pain*, *Neg_Feeling*, *Wish*, and *CI_Regret*. The following three examples illustrate the design of the situations we ask our participants to imagine (*Neg_Feeling*, *Body1*, and *Feeling*)⁷:

⁷The conversations for all the stimuli as well as training material on what it means to be a contradiction can be found here: https://osf.io/8wuh6?view_only=6f64b8528d08495c824ed6f53ecc9dcd.

**Neg_Feeling
(Regret)**

Tom says to Sally: 'I regret drinking instant coffee this morning.'
 Sally responds: 'Oh, so you mean that you have a negative feeling about drinking instant coffee this morning.'
 Tom responds: 'No, I don't mean to say that. I have a positive feeling about drinking instant coffee this morning.'

Body₁ (Pain)

Tom says to Sally: 'I have a pain in my arm.'
 Sally responds: 'Oh, so you mean that there is something physically wrong with your arm?'
 Tom responds: 'No, I don't mean to say that. My arm is perfectly fine.'

Feeling (Pain)

Tom says to Sally: 'I have a pain in my arm.'
 Sally responds: 'Oh, so you mean that you're feeling something unpleasant in your arm?'
 Tom responds: 'No, I don't mean to say that. My arm feels perfectly fine.'

Please note that we use the same labels in Study 2 and in Study 1 as we tested for the same content, albeit using different stimuli. After being presented with the conversation, participants were then asked the question, 'Does Tom contradict himself?' The participants answered using a 9-point Likert scale anchored at '1 = definitely not' and '9 = definitely yes'.

V.2 Hypotheses

As preregistered (<https://osf.io/kqnc8>), we investigated the following hypotheses:

- H₄:** The content of *Body₁*, *Body₂*, and *Feeling*, as well as the content of *Neg_Feeling* and *Wish*, receive contradiction ratings that are significantly above the midpoint of 5.
H₅: The contents of *CI_Pain* and *CI_Regret* receive contradiction ratings that are significantly below the midpoint of 5.

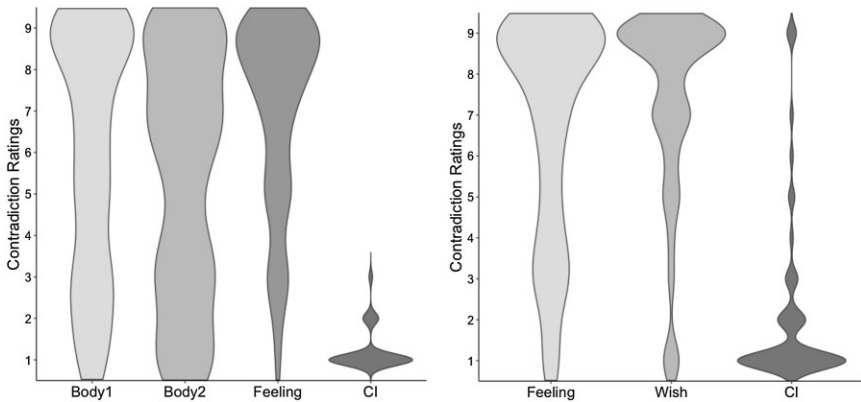
Before the tests, the participants were given a short description of what it means for speakers to *contradict* themselves in the philosophically relevant sense. They then answered two test questions that served as comprehension checks. We excluded two participants who failed both of these test questions. The remaining 408 participants had a mean age of 38.75 years, with 107 indicating 'male', 295 'female', and six 'non-binary' as their gender.

V.3 Results

The mean ratings and statistical results for each of the seven conversations are listed in Table 4 and illustrated in Fig. 2. We conducted *t*-tests to identify those

Table 4. Mean contradiction ratings for pain and regret conditions and results of *t*-tests against the midpoint of the scale.

Concept	Condition	Mean	Standard error	<i>t</i>	<i>p</i> -value
Pain	Body1	5.95	0.395	2.405	=0.010
	Body2	5.56	0.383	1.460	=0.075
	Feeling	7.27	0.286	7.945	<0.001
	CI_Pain	1.22	0.065	-57.846	<0.001
Regret	Neg_Feeling	6.49	0.354	4.214	<0.001
	Wish	7.74	0.287	9.534	<0.001
	CI_Regret	2.11	0.280	-10.357	<0.001

**Figure 2.** Violin plots for the four pain conditions (left) and the three regret conditions (right) showing the distribution of the contradiction ratings.

conversations for which the contradiction ratings were significantly above the midpoint of 5. Except for *Body2*, **H4** was supported for all conditions for which we expected high contradiction ratings. Rating for both conversational implicatures was significantly below the midpoint of 5, providing evidence for **H5**. Subsequent pairwise comparisons indicated no significant difference between *Body1* and *Body2* ($p = 1.00$), but both bodily conditions were significantly different from the *Feeling* condition ($p = 0.026$ and $p = 0.002$).⁸

⁸ We conducted an ANOVA to compare the three main conditions (*Body1*, *Body2*, and *Feeling*) in the pain condition. The independent factor was found to be statistically significant ($F(2,173) = 50.04$, $p = 0.002$, $\eta^2 = 0.07$).

V.4 Discussion

In the pain condition, *Body1* and *Feeling* received ratings that were significantly above the midpoint, thus suggesting that both contents were semantically entailed by the claim ‘I have a pain in my arm.’ While we do not have a completely satisfactory explanation for why the ratings for *Body2* were slightly decreased, it is possible that the precise wording of the condition had a negative effect on people’s ratings. Section VI.2 provides a more detailed discussion of the differences in ratings between *Body1* and *Body2*, as well as between both bodily conditions and *Feeling*. Finally, *CI_Pain* received low contradiction ratings, thus indicating that its content was only conversationally implicated. The regret condition functioned as expected. The contents of *Neg_Feeling* and *Wish* were considered to be semantically entailed by the regret statement, whereas the content of *CI_Regret* appeared to be only conversationally implicated given the low contradiction ratings for the respective conversation.

VI. GENERAL DISCUSSION

VI.1 Summary of the results

Our investigation aimed to better understand the folk concept of pain using a novel methodological approach based on pain linguistics. One of the central questions in the philosophy of pain is what the semantic features of the folk concept of pain are. Three suggestions have been made in the literature: The first is that the semantic content mainly pertains to a feeling (feeling view), the second is that it mainly pertains to a bodily state (bodily view), and the third is that it includes both feeling and bodily information (pluralist view). To distinguish among these three options experimentally, we focused on first-person pain reports using the implication, projection, and deniability tests.

The results of Study 1 and Study 2 indicate that both information about a bodily state and an unpleasant feeling meet our criteria for semantic entailment, as depicted in Table 5. First, the information that there is something wrong with Tom’s arm, that Tom thinks that there is something wrong with his arm, and that Tom has an unpleasant feeling was reliably inferred from Tom’s statement, ‘I have a pain in my arm.’ Secondly, these three implications did not project when embedded in an entailment-cancelling operator: Most participants no longer inferred a bodily or feeling component when Tom said, ‘I don’t have a pain in my arm.’ Thirdly, the target contents cannot be denied when a first-person pain report is made without producing a contradiction. At the same time, the intuitively unrelated, presupposed, and conversationally implicated contents behaved as expected in relation to the target statement ‘I have a pain in my arm.’

Table 5. Results of the implication, projection, and deniability test for (potential) implications of the statement ‘I have a pain in my arm.’

	Implication	Projection	Deniability
Body1	☑	☒	☒
Body2	☑	☒	☒
Feeling	☑	☒	☒
Unrelated	☒		
Presup_Pain	☑	☑	
CI_Pain	☑	☒	☑

These results allow us to draw a set of conclusions about our research, which we will discuss in the following. To do so, we first offer a set of methodological remarks (Section VI.2). Secondly, we locate our results within a broader perspective to show that they favour a pluralist view while acknowledging some tensions with vignette-based and corpus-linguistic studies (Section VI.3).

VI.2 Methodological limitations

Our investigations reveal that well-established linguistic tests, namely the implication, projection, and deniability tests, prove to be useful in their application to first-person pain reports. The regret condition that served as our control confirmed that the experiments were well-designed. We consider this to be an innovative shift in the ongoing philosophical discourse, one that has the potential to uncover previously overlooked facets of the common understanding of pain. Therefore, pain linguistics can provide new insights into ordinary pain language. While we consider our experimental framework to be promising, we would also like to discuss four potential limitations.

First, our methodological approach allows us to avoid the (lack of) competence and performance effects that may alter people’s responses in vignette-based studies. However, first-person pain reports may also provoke certain biases. While the presentation of first-person pain reports was relatively context-free in our studies, the participants may have already had a particular context in mind, in which the corresponding statements typically occur. Our design cannot reveal the various contexts that the participants might have imagined when reading the target statement, ‘I have a pain in my arm.’ Nonetheless, in the absence of further contextual information, we assume that, in general, our experimental framework triggered thoughts about a broad array of conversational situations we aimed to test.

Secondly, one could potentially raise concerns regarding the efficacy of the proposed experimental design, and whether it poses a sufficient level of rigour. There may be a risk that our semantic entailment criteria could be met

too readily, thereby compromising the validity of the study. For example, one might question whether a rating above the midpoint in the implication test is sufficient to establish that content is reliably inferred. Similarly, one might question whether a rating above the midpoint in the deniability test is sufficient to establish that a target content cannot be denied without producing a contradiction. Previous studies that have used the cancellability or deniability tests (see e.g. Willemsen and Reuter 2021 on thick concepts) have revealed that average results between 6 and 7 were highly common for semantically entailed content and thus reliably indicated such a semantic component. Results between 7 and 9 are rare, presumably because (a) several participants shy away from the endpoints of a scale, (b) the statements that are used in cancellability and deniability tasks are not always easy to comprehend, and (c) some participants might entertain unusual interpretations of the terms in question.

Thirdly, although we might be justified in claiming that both bodily and feeling aspects are semantically entailed components, the differences in the ratings for *Body1* (5.95), *Body2* (5.56), and *Feeling* (7.27) in the deniability test require further explanation—especially as the pairwise comparison indicated no significant difference between *Body1* and *Body2*, but between both bodily conditions and *Feeling*. One potential reason is the particular wording of the bodily conditions. To say that there is something physically wrong with a body part is likely to communicate a certain type of severity that not everyone associates with a damaged, disturbed, or disrupted body state (Liu 2020). For example, some participants might understand Tom's statement 'I have a pain in my arm' as referring to a bodily state that is not considered to be as serious as the expression 'physically wrong' would suggest, such as in the case of a slight bruising or overextension of the arm. Similarly, not all pains indicate 'wrongness', for instance, having sore muscles after exercising might actually be considered exactly how one's body *should* feel. Unfortunately, there does not seem to be a common term that comprises all those various phenomena like bruises, injuries, sore muscles, etc. We opted for the phrasing 'physically wrong' due to its inherent generality, encompassing a range of bodily disorders while also indicating a departure from the normal state of the affected body part.

There may be an additional reason for the lower ratings of the bodily conditions compared to the feeling condition, namely the possibility that respondents were contemplating cases of referred pain. In such cases, patients experiencing a heart attack or spinal disc herniation may report pain in their arms, despite the fact that their arms are structurally sound. That is, while the person subjectively experiences pain in one body part, the corresponding physiological disturbance is located in another body part. Following an objective interpretation of the bodily view, participants who are aware of such referred pain cases are likely *not* to consider the two statements 'I have a pain

in my arm' and 'My *arm* is perfectly fine' to be contradictory. They should, however, consider the statements 'I have a pain in my arm' and 'My *body* is perfectly fine' to be contradictory. We chose to be more specific about the bodily location, as it represents a more natural way of speaking. However, this might have led some participants to give rather low contradiction ratings.

Finally, we would like to address the question regarding the compatibility of the results of vignette-based research and our experiments. As stated above, vignette-based studies have yielded results that seem to support unitary views. It is important to note that many vignette studies were designed in such a way that feeling and bodily views were pitted against each other. Thus, the participants' responses would either speak against the bodily view or against the feeling view. Such studies can and do reveal important insights into people's thoughts about pain but, as we have argued above, they might also predispose people to provide certain answers due to the frequently unusual contexts and designs. Thus, this might also reveal that different factors activate different conceptions in our thinking of pain (Borg *et al.* 2021). One of the most obvious factors is, of course, the first-person communicative act that was at the centre of our investigation. Such a communicative act might introduce confounding factors such as charitable interpretations on part of the listener for which we were not able to control. Examining such factors was beyond the scope of our paper but should be included in future studies.

Overall, we need to remain cautious about whether our three tests can jointly prove what constitutes part of the semantic content of the folk concept of pain. In our opinion, this does not generally speak against our experimental framework but underlines the complexity of the folk concept of pain and the methodological ingenuity needed for its investigation. As indicated in Section II, we need a new methodological approach to the folk concept of pain that captures the linguistic intuitions of participants in as natural a context as possible. Focusing on pain linguistics enables us to go beyond previous vignette-based studies. Naturally, this approach also has certain methodological limitations. Therefore, we consider vignette studies and our experimental-linguistic design to complement one another.

VI.3 *Theoretical implications*

The findings from our studies may appear to offer favourable prospects for both the bodily and feeling views, given that bodily and feeling components were found to be semantically entailed in first-person pain expressions. However, both views seem to tell only half of the story. By contrast, the pluralist view is more commensurate with the results of our studies. Information regarding both a bodily disruption and an unpleasant feeling appears to be communicated as part of the semantic content of the folk concept of pain in first-person pain reports. This suggests that the folk concept of pain is indeed complex. That is, the paradox of pain cannot be solved by denying that one of the two features

is part of the semantic content of the folk concept of pain, as suggested by unitary views. Instead, a pluralist view acknowledges the complexity of the semantic features, thus rejecting the assumption of a (single) folk concept of pain with a univocal meaning.

As indicated in Section II, pluralist accounts take multiple forms. In particular, they differ in how they characterize the relationship between feeling and bodily features as part of the semantic content of the folk concept of pain. It is likely that not all pluralistic approaches will be compatible with the combined results of previous vignette-based and our experimental-linguistic studies. Therefore, the different pluralistic approaches deserve closer examination focusing on the polyeidic, polysemous, and component views as the most prominent pluralistic approaches defended in the recent literature.

Borg, Hansen & Salomons (2019) and Borg *et al.* (2020) defend a *polyeidic view*, which posits a single folk concept of pain that is understood as an amalgam of multiple dimensions. One of these dimensions has mental and bodily aspects as its opposite end points. The polyeidic view allows some people to tend towards the mental end point of the dimension and others to tend towards the bodily end point. Furthermore, the polyeidic view allows the same person to treat pain as a mental state in one context and as a bodily state in another. Central to the polyeidic view is its dimensional character: Based on contextual and individual differences, people may tend towards one or the other direction of the spectrum (Borg *et al.* 2021). However, according to the polyeidic view, the same person cannot treat pain in the same situation as an unpleasant feeling and a disruptive bodily state:

To return to the question of whether the Polyeidic view is committed to maintaining that the folk view of pain is paradoxical: this version of the view would seem to allow that it was not, since no experience of pain would ever be conceptualized as, for example, both entirely mental or entirely bodily at one and the same time (Borg *et al.* 2020: 44).

The polyeidic view encounters two challenges. First, the distinction of treating pain as a bodily or mental state appears to be categorical rather than dimensional; this is in contrast to some of the other properties included in the amalgam of dimensions (for example, sensory, affective, and motivational), which may well be modelled as having different degrees (Coninx 2022). Secondly, and most relevant for the present purposes, the polyeidic view appears to contradict the results of our implication test. The same participants infer feeling and bodily information from first-person pain reports at the same time. To account for these data, feeling and bodily features cannot be modelled as located at the opposite ends of the same dimension. Instead, we need an account that predicts that people will treat pain as simultaneously involving feeling and bodily aspects, at least in ordinary cases.

Liu (2021, 2022) defends a *polysemy view*. Most views that have been defended in the recent literature presuppose that there is only a single folk concept of pain that has either a univocal or a more complex meaning. By contrast, the

polysemy view suggests that there are two distinct folk concepts of pain with univocal, although related, meanings: One that treats pain as a mental state and one that treats pain as a bodily state. Thus, the polysemy view assumes that there are two related folk concepts of pain. Interestingly, in many prototypical cases of polysemy (e.g. line), only one of the related concepts is employed with sufficient context (e.g. drawing a line vs reading a line). Instead, the apparently separate concepts associated with the term ‘pain’, referring to a mental state and a bodily state, respectively, are systematically derived together as indicated by our experiments. Thus, the polysemy view is compatible with the empirical findings provided thus far only if it can plausibly allow for the two postulated concepts to be employed simultaneously.

Thirdly, the *component view* presupposes a single folk concept of pain that is composed of multiple elements: The folk concept of pain includes both feeling as well as bodily aspects. In principle, there are strong and weak versions of the component view. Defenders of a strong version assume that information about a disruptive bodily state and an unpleasant feeling are necessary components of the semantic content. In this version, the results of existing vignette studies remain to be explained. By contrast, advocates of a weaker version of the component view (e.g. Corns 2020) may admit that people also ascribe pain in the absence of one of these two components in exceptional circumstances, even if these cases are likely to be perceived as less paradigmatic. Thus, this approach reveals strong ties to family resemblance (Wittgenstein 1958) or prototype theories (Rosch and Mervis 1975). The difficulty at hand is to thoroughly examine the conditions under which an attribution of pain still takes place and to determine which cases are considered more or less pragmatic.

VII. CONCLUSION

In this paper, we argue for a pluralistic approach to the folk concept of pain in a two-fold manner. First, pain linguistics provides a new methodological approach to the folk concept of pain that not only circumvents some of the limitations of vignette studies but also does justice to the complexity and context-sensitivity of folk concepts. Secondly, the results of our studies allow us to identify pluralist accounts of pain as being more promising than unitary views. The three tests we conducted indicate that feeling and bodily content are *both* semantically entailed by first-person pain reports. Our studies make yet another theoretical contribution, as only the polysemy and component views are compatible with our results. In particular, the component view deserves further investigation, as it has not yet been sufficiently considered in recent debates.⁹

⁹ We would like to express our gratitude to Lucien Baumgartner, Severin Keller, Michelle Liu, Justin Sytsma as well as the participants of the Zurich XPhi Lab Meeting for their helpful feedback. Furthermore, we would like to thank Albert Newen for joint discussions that initiated the project.

FUNDING

The research of Sabrina Coninx was funded by the German Research Foundation (DFG), project number GRK-2185/1 (Research Training Group Situated Cognition). The research of Kevin Reuter was funded by the Swiss National Science Foundation (SNSF), grant number PCEFP1_181082. Pascale Willemsen also received generous support from the SNSF, grant numbers PCEFP1_181082 and PZooP1_201737.

REFERENCES

- Almeida, G., Struchiner, N. and Hannikainen, I. (2023) 'Rule Is a Dual Character Concept', *Cognition*, 230: 105259.
- Austin, J. L. (1962) *How To Do Things with Words*, Oxford: OUP.
- Aydede, M. (2009) Pain E. N. Zalta (ed.) spring 2013 edn. <<http://plato.stanford.edu/archives/spr2013/entries/pain/>> accessed 29 May 2009.
- Bach, K. (2006) 'The Top 10 Misconceptions about Implicature', in B. J. Birner and G. Ward (eds) *Drawing the Boundaries of Meaning*, 21–30. Amsterdam: John Benjamins Publishing Company.
- Baumgartner, L., Willemsen, P. and Reuter, K. (2022) 'The Polarity Effect of Evaluative Language', *Proceedings of the Annual Meeting of the Cognitive Science Society*, 44: 729–35.
- Blome-Tillmann, M. (2008) 'Conversational Implicature and the Cancellability Test', *Analysis*, 68: 156–60.
- Borg, E., Hansen, N. and Salomons, T. (2019) 'The Meaning of Pain Expressions and Pain Communication', in S. van Rysewyk (ed.), *Meanings of Pain*, 2nd edn. 261–82. Cham: Springer.
- Borg, E. et al. (2020) 'Is the Folk Concept of Pain Polyeidic?', *Mind & Language*, 35: 29–47.
- . et al. (2021) 'Pain Priors, Polyeidism, and Predictive Power: A Preliminary Investigation into Individual Differences in Ordinary Thought about Pain', *Theoretical Medicine and Bioethics*, 42: 113–35.
- Chierchia, G. and McConnell-Ginet, S. (2000) *Meaning and Grammar. An Introduction to Semantics*, 2nd edn. Cambridge: MIT Press.
- Coninx, S. (2022) 'A Multidimensional Phenomenal Space for Pain: Structure, Primitiveness, and Utility' *Phenomenology and the Cognitive Sciences*, 21: 223–43.
- . (2020) *Experiencing Pain: A Scientific Enigma and Its Philosophical Solution*. Berlin, Boston: De Gruyter.
- Corns, J. (2020) *The Complex Reality of Pain*. New York: Routledge.
- Davis, W. (2019) *Implicature* E. Zalta (ed.) fall 2019 edn. <<https://plato.stanford.edu/archives/fall2019/entries/implicature/>> accessed 6 Sep 2019).
- Goldberg, B., Reuter, K. and Sytma, J. (2023) 'The History of the Concept of Pain: How the Experts Came To be out of Touch with the Folk', in K., Hens and A. De Block (eds) *Advances in Experimental Philosophy of Medicine*. London: Bloomsbury Academic.
- Grice, H. P. (1975) 'Logic and Conversation', in P. Cole and J. L. Morgan (eds) *Syntax and Semantics 3: Speech Acts*, 41–58. New York: Academic Press.
- Hill, C. S. (2005) 'Ow! the Paradox of Pain' in M. Aydede (ed.) *Pain: New Essays on Its Nature and the Methodology of Its Study*, 75–98. Cambridge: MIT Press.
- Horn, L. R. (1986) 'Presupposition, Theme and Variations', in A. M. Farley, P. T. Farley and K.-E. McCullough (eds) *Papers from the Parasession on Pragmatics and Grammatical Theory at the Twenty-Second Regional Meeting*, 168–92. Chicago: Chicago Linguistic Society.
- Huang, Y. (2006) 'Neo-Gricean Pragmatics', in K. Brown (ed.), *Encyclopedia of Language & Linguistics*, 2nd edn. 586–90. Amsterdam: Elsevier.
- Hyman, J. (2003) 'Pains and Places', *Philosophy*, 78: 5–24.
- Kim, H. et al. (2016) 'Where Is Your Pain? A Cross-cultural Comparison of the Concept of Pain in Americans and South Koreans', *Studia Philosophica Estonica*, 9: 136–69.

- Kneer, M. (2021). 'Predicates of Personal Taste: Empirical Data', *Synthese*, 199: 6455–71.
- . (2022). 'Contextualism Versus Relativism: More Empirical Data', in J. Wyatt, J. Zakkou and D. Zeman (eds) *Perspectives on Taste: Aesthetics, Language, Metaphysics, and Experimental Philosophy*, 109–40. New York: Routledge.
- Kripke, S. A. (1981) *Naming and Necessity*, Cambridge, MA: Harvard University Press.
- Levinson, S. C. (1983) *Pragmatics*. Cambridge: CUP.
- Lewis, D. (1980) 'Mad Pain and Martian Pain', in N. Block (ed.) *Readings Philosophy of Psychology*, vol. 1, 216–22. Cambridge, MA: Harvard University Press.
- Liu, M. (2020) 'The Intuitive Invalidity of the Pain-in-Mouth Argument', *Analysis*, 80: 463–74.
- . (2021) 'Pain, Paradox, and Polysemy', *Analysis*, 81: 461–70.
- . (2022) 'The Polysemy View of Pain', *Mind & Language*, 38: 198–217.
- Machery, E. (2017) *Philosophy within Its Proper Bounds*. Oxford: OUP.
- Massin, O. (2017) 'Bad by Nature: An Axiological Theory of Pain', in J. Corns (ed.) *The Routledge Handbook of Philosophy of Pain*, 321–33. New York: Routledge.
- Muth, C., Briesen, J. and Carbon, C.-C. (2020) "'I Like How It Looks but It Is Not Beautiful": Sensory Appeal beyond Beauty', *Poetics*, 79: 101376.
- Raja, S. N. et al. (2020) 'The Revised International Association for the Study of Pain Definition of Pain: Concepts, Challenges, and Compromises', *Pain*, 161: 1976–82.
- Reins, L. M. and Wiegmann, A. (2021) 'Is Lying Bound to Commitment? Empirically Investigating Deceptive Presuppositions, Implicatures, and Actions', *Cognitive Science*, 45: e12936.
- Reuter, K. (2011) 'Distinguishing the Appearance from the Reality of Pain', *Journal of Consciousness Studies*, 18: 94–109.
- . (2017) 'The Developmental Challenge to the Paradox of Pain' *Erkenn*, 82: 265–83.
- Reuter, K., Phillips, D. and Sytsma, J. (2014) 'Hallucinating Pain', in J. Sytsma (ed.) *Advances in Experimental Philosophy of Mind*, 75–100. London: Bloomsbury Academic.
- Reuter, K., Sienhold, M. and Sytsma, J. (2019) 'Putting Pain in Its Proper Place', *Analysis*, 1: 72–82.
- Reuter, K. and Sytsma, J. (2020) 'Unfelt Pain', *Synthese*, 197: 1777–801.
- Rosch, E. and Mervis, C. B. (1975) 'Family Resemblances: Studies in the Internal Structure of Categories', *Cognitive Psychology*, 7: 573–605.
- Salomons, T. V. et al. (2022) 'Is Pain "All in Your Mind"? Examining the General Public's Views of Pain', *Review of Philosophy and Psychology*, 13: 683–98.
- Sytsma, J. (2010) 'Dennett's Theory of the Folk Theory of Consciousness', *Journal of Consciousness Studies*, 17: 107–30.
- Sytsma, J. and Machery, E. (2009) 'How To Study Folk Intuitions about Phenomenal Consciousness', *Philosophical Psychology*, 22: 21–35.
- Sytsma, J. and Reuter, K. (2017) 'Experimental Philosophy of Pain', *Journal of Indian Council of Philosophical Research*, 34: 611–28.
- Sytsma, J., Willemsen, P. and Reuter, K. (2022) 'Mutual Entailment between Causation and Responsibility'. (Pre-print: <http://philsci-archival.pitt.edu/20497/>)
- Tye, M. (2017) 'Are Pains Feelings?' *The Monist*, 100: 478–84.
- Väyrynen, P. (2013) *The Lewd, the Rude and the Nasty*. Oxford: OUP.
- Willemsen, P. and Reuter, K. (2021) 'Separating the Evaluative from the Descriptive: An Empirical Study of Thick Concepts', *Thought*, 10: 135–46.
- Wittgenstein, L. (1958) *Philosophical Investigations*, 2nd edn. (G. E. M. Anscombe (trans.)). Oxford: Blackwell.
- Zakkou, J. (2018) 'The Cancellability Test for Conversational Implicatures', *Philosophy Compass*, 13: e12552.

¹ *Vrije Universiteit Amsterdam, The Netherlands*

² *University of Zurich, Switzerland*