

- Lewis, D. 1980. A subjectivist's guide to objective chance. In *Studies in Inductive Logic and Probability, Volume II*, ed. R.C. Jeffrey, 263–93. Berkeley: University of California Press.
- Lewis, D. 1986. *On the Plurality of Worlds*. Oxford: Blackwell.
- Pickel, B. and M. Schulz 2018. Quinean updates: in defense of “Two dogmas”. *Journal of Philosophy* 115: 57–91.
- Quine, W.V.O. 1951. Two dogmas of empiricism. *Philosophical Review* 60: 20–43.
- Skyrms, B. 1995. Strict coherence, sigma coherence and the metaphysics of quantity. *Philosophical Studies* 77: 39–55.
- Sprenger, J. 2018. The objectivity of subjective Bayesianism. *European Journal of Philosophy of Science* 8: 539–58.
- White, R. 2013. Evidence cannot be permissive. In *Contemporary Debates in Epistemology*, 2nd edn, eds. M. Steup, J. Turri, and E. Sosa, 312–23. New York: Blackwell.

## *Pain, paradox and polysemy*

MICHELLE LIU 

### 1. Introduction

The term ‘paradox of pain’ refers to the claim that the folk concept of pain is paradoxical, treating pains as simultaneously *mental* and *bodily* (Hill 2005, 2017, Aydede 2013, Reuter et al. 2014, Reuter 2017, Borg et al. 2020). Consider the characterization of the paradox of pain in a recent paper by Borg et al.:

Pain states . . . seem to be classic Cartesian mental states, possessing the characteristic hallmarks of such states: subjectivity, privacy, and incorrigibility. Furthermore, philosophers have assumed that this sort of mentalised view captures the ordinary *folk* notion of pain. . . . However, a moment's reflection on how people ordinarily think and talk about pain shows that the folk view may in fact be somewhat more complex than this model suggests, for folk clearly hold that pains, unlike other paradigm mental states like beliefs, have (non-brain-based) bodily locations. . . . Recognition of this fact has led some theorists (e.g. Hill 2005) to talk of the ‘paradox of pain’ . . . (2020: 29–30)

The passage identifies two conceptions of pain. On the one hand, pain seems to be a *mental state*, supported by the claim that pain possesses characteristics that conscious mental states are often said to have. These characteristics, as mentioned by Borg et al., are (see also Hill 2005, Aydede 2013, Reuter et al. 2014, Hardcastle 2015: 531–32, Reuter 2017):

Privacy: One has privileged access to one's own pains.

Subjectivity: If one has pain, then one feels pain.

Incorrigibility: If one feels pain, then one has pain.

Subjectivity and incorrigibility jointly entail that there is no appearance–reality distinction with respect to pain, that is, there is no difference between *pain* and the *experience of pain* (Kripke 1980: 154). Indeed, philosophers of mind have traditionally treated pain as an archetypical example of phenomenal experience (e.g. Lewis 1980, Tye 2017). We call the conception of pain that treats pain as a mental state the ‘mental conception of pain’.

On the other hand, the folk view of pain also suggests a ‘bodily conception of pain’, treating pains as states of the body, located in non-brain-based body parts. As Borg et al. (2020: 30) note, ‘When someone stubs their toe or cuts their hand the pain is (or at least is said to be) *in the toe* or *in the hand*.’

However, body parts are not the type of entities that can have mental states. In the literature, the tension is often put in terms of *location*: ‘mental states are not the kind of entities that inhabit [non-brain-based] body parts’ (Reuter 2017: 265); they are ‘in the head [i.e. brain], if they are anywhere’ (Aydede 2013). If the folk concept of pain treats pain as simultaneously a mental state and a bodily state, it would seem that the folk concept is paradoxical.<sup>1</sup>

Theorists have given different responses to the alleged paradox of pain. Hill (2005, 2017), who coined the term ‘the paradox of pain’, suggests that the folk concept of pain should be revised. Hill thinks that we should introduce two separate concepts: ‘peripheral pain’, which coheres with the bodily conception, and ‘central state pain’, which coheres with the mental conception. Reuter (2017), Reuter et al. (2014) and Reuter and Sytsma (2020) deny the alleged paradox. Appealing to empirical data, they argue that people have only a bodily conception of pain, not a mental conception. Borg et al. (2020) argue that the evidence presented by Reuter et al. does not demonstrate that people lack a mental conception of pain. They think that the ordinary concept of pain is ‘polyeidic’, ‘containing different strands or elements’, including both mental and bodily elements, where ‘in different contexts different elements of the concept could be activated’ (2020: 30–31). On this view, while a clinician might treat pain as a bodily state and a patient might treat pain as a mental state, they still ‘share a common concept of pain’ (2020: 45). Proponents of this view take the polyeidic nature of the concept of pain to accommodate the intuition that the folk view of pain is paradoxical without entailing that it is paradoxical. They claim that while one might take pain to be

1 Hill’s (2005, 2017) way of spelling out the paradox of pain is more nuanced. For Hill, the referent of our folk concept of pain is tissue damage, but the conceptual role that governs the application of our concept of pain suggests pain as perceptual representation (i.e. mental state) rather than what is being represented (i.e. tissue damage).

It is also worth noting that some theorists, instead of speaking of the folk concept of pain as paradoxical, take the aforementioned two conceptions of pain to be in tension in the sense that they lead to a strange view of pain which treats pains as located in public space (i.e. body parts), while being private, subjective and incorrigible (see e.g. Hardcastle 2015: 531).

mental in one context and bodily in another, no one would take it to be both mental and bodily at the same time (2020: 44).

All of these proposals agree that there is a single folk concept of pain. The question is whether this concept is internally coherent. This paper questions the consensus that there is a single concept of pain. It argues that our primary pain terms – ‘pain’, ‘sore’, ‘ache’ and ‘hurt’ – are polysemous (§2). This polysemy reflects two concepts of pain, and once this is recognized the paradox of pain is resolved (§3). The paper concludes with wider implications for the philosophical debate about the nature of pain (§4).

## 2. Pain terms and polysemy

In English, there are four *primary pain terms* – ‘pain’, ‘sore’, ‘ache’ and ‘hurt’ – which are words specifically dedicated to denoting pain (Fabrega and Tyma 1976, Reznikova et al. 2012).<sup>2</sup> I will discuss the last three pain terms first, because they behave in a similar way syntactically, and then turn to ‘pain’.

Consider typical pain reports using the three pain terms – ‘sore’, ‘ache’ and ‘hurt’:

- (1) a. My leg is sore.
- b. My arm is aching.
- c. My thumb is hurting.

In (1)-sentences, a body part is described as being in a particular state. On the level of surface grammar, pain is treated as a property of a body part, indicated by either an adjectival phrase, for example ‘sore’ in (1-a), or a verb, for example ‘ache’ and ‘hurt’ in their respective progressive forms in (1-b) and (1-c). Note that the same predicates also apply to persons:

- (2) a. I am sore.
- b. I am aching.
- c. I am hurting.

In (2)-sentences, the surface grammar describes the person or the experiencer, rather than the body part, as being in a particular state. I argue that these predicates, which apply to both body parts and persons, are polysemous.

*Polysemy* is a linguistic phenomenon where a word has multiple related meanings or senses. Polysemy is contrasted with *homonymy*, where a word form has distinct but unrelated meanings, for example ‘bank’ (financial institution vs. river edge). Instances of polysemy and homonymy fall under the phenomenon of *ambiguity*. An expression is *ambiguous* if it has multiple

2 *Primary pain terms* are contrasted with *secondary pain terms* such as ‘burning’, ‘stinging’, ‘shooting’, which are not exclusively tied to pain reports.

meanings, where the meanings are either related (in case of polysemy) or unrelated (in the case of homonymy).

A standard test for ambiguity is the Conjunction Reduction test (Sennet 2016; see also Chomsky 1957: 35–36). Sentence (3-c) is the conjunction reduction of (3-a) and (3-b):

- (3) a. Tom is tired.
- b. Jerry is tired.
- c. Tom and Jerry are tired.

Sentence (3-c) makes perfect sense.

Now consider the polysemous word ‘expire’, which means ‘die’ in (4) and ‘become invalid’ in (5):

- (4) Tom expired.
- (5) Tom’s student card expired.

The conjunction reduction of (4) and (5) sounds odd:

- (6) \*Tom and his student card expired.

Sentence (6) is an example of what is known as ‘zeugma’: it is not outright nonsensical, but it sounds odd because it is difficult to come up with a straightforward interpretation. There is no single property of *being expired* which can be attributed to both Tom and his student card.

To test whether an expression is ambiguous, we can conjoin phrases within the scope of the expression to see if the resulting sentence is zeugmatic. If zeugma results, then the expression is ambiguous – it is either polysemous or homonymous.

Let us now apply the Conjunction Reduction test to the primary pain terms ‘sore’, ‘ache’ and ‘hurt’. First consider the following sentences where the subject is a conjunction of either body parts or persons but not both:

- (7) a. John’s arms and my legs are sore.
- b. Tom and Lisa are aching.
- c. Harry’s toes and my fingers are hurting.

These (7)-sentences sound fine. However, the following sound infelicitous:

- (8) a. \*John’s arms and I are sore.
- b. \*Tom’s hands and Lisa are aching.
- c. \*Harry and my fingers are hurting.<sup>3</sup>

3 The linguistic intuitions here have been empirically confirmed: 107 native English speakers on Amazon Mechanical Turk (MTurk) were asked whether (7)-sentences and (8)-sentences sounded odd on a 7-point Likert scale where ‘1’ means ‘This sentence definitely does not sound odd’ and ‘7’ means ‘This sentence definitely sounds odd’. ANOVA yielded no

In (8)-sentences, the subject is a conjunction of a noun phrase (NP) indicating body parts with another NP indicating a person. These (8)-sentences are zeugmatic because there is no single property of *being sore* or *aching* or *hurting* that applies to both NPs in the relevant subject position. The two NPs require the respective predicate to express different properties.

Another test for ambiguity is the Ellipsis test, which works in a similar manner to the Conjunction Reduction test (Lakoff 1970, Sennet 2016). Consider the following elliptical construction:

- (9) Larry is tired and so is Jerry.

Sentence (9) has a cross-interpretation – the elliptical phrase ‘so is Jerry’ is understood to mean ‘Jerry is tired’ which is the elided phrase. Now consider (10), which contains the polysemous word ‘expire’:

- (10) Tom expired and so did his card.

Sentence (10) is odd or zeugmatic. With a cross-interpretation, ‘so did his card’ would have to mean ‘his card expired (died)’. A card can be said to ‘expire’, but not in the sense that Tom can be said to ‘expire’. The predicate ‘expire’ applies to persons such as Tom and inanimate objects such as cards in different senses.

To test for ambiguity, we can construct an ellipsis to see whether a cross-interpretation results in zeugma.

Now consider primary pain terms again. The following elliptic constructions allow cross-interpretations and sound fine:

- (11) a. John’s arms are sore and so are my legs.  
 b. Tom is aching and so is Lisa.  
 c. Harry’s toes are hurting and so are my fingers.

In contrast, the following are zeugmatic:

- (12) a. \*John’s arms are sore and so are we.  
 b. \*Tom’s hand is aching and so is Lisa.  
 c. \*Harry is hurting and so is my thumb.<sup>4</sup>

---

significant difference among the three sentences in each group ((7)-sentences:  $F(2, 249) = 0.47, p = 0.62$ ; (8)-sentences:  $F(2, 249) = 0.10, p = 0.90$ ). Average ratings for (7)-sentences and (8)-sentences were 2.75 (SD = 2.07) and 6.03 (SD = 1.54) respectively. A paired-sample t-test of the two groups showed significant differences between average ratings of (7)-sentences and (8)-sentences ( $t(83) = 19.84, p < 0.001$ ).

4 These linguistic intuitions have been empirically confirmed. Results from 76 native English speakers showed average ratings for (11)- and (12)-sentences of 2.15 (SD = 1.63) and 4.88 (SD = 1.89) respectively. A paired-sample t-test showed significant differences between the two groups ( $t(67) = 13.64, p < 0.001$ ). Average ratings for each of the (11)- and (12)-sentences were: 1.82 (SD = 1.33), 2.47 (SD = 1.90), 2.16 (SD = 1.66), 5.24 (SD = 1.83), 5.16 (SD = 1.82), 4.25 (SD = 1.89). Pairwise comparisons showed that (12-c) was rated as less odd than (12-a) ( $t(67) = 4.28, p < 0.001$ ) and (12-b) ( $t(67) = 3.44, p < 0.001$ ).

These (12)-sentences do not allow cross-interpretations without sounding somewhat odd.

The failure of both the Conjunction Reduction test and the Ellipsis test demonstrates that primary pain terms such as ‘sore’, ‘ache’ and ‘hurt’ are ambiguous – they apply to a person and a person’s body part in different senses. If they are ambiguous, then they are either polysemous (where the meanings are related) or homonymous (where the meanings are unrelated). The hypothesis that they are homonyms is quickly ruled out. Clearly the relevant meanings are related – if S’s body part is described as sore/aching/hurting, then S can also be described as sore/aching/hurting. So these pain terms must be *polysemous*.

A case for polysemy can also be made for the primary pain term ‘pain’ itself. A common way to report pains in English is to use the phrase ‘in pain’:

(13) I am in pain.

Sentence (13) is used to report physical pain, for example when someone is injured, as well as emotional pain, for example when someone is grieving. In these cases, pain is treated as a mental state of a person or a sentient being. The phrase ‘in pain’ does not apply to body parts:

(14) \*My leg is in pain.

Sentence (14) is odd because body parts are not the kind of entities that are capable of having mental states. Nevertheless, we also speak of pains of body parts, modifying the noun ‘pain’ with a prepositional phrase indicating a body part:

(15) There is a pain in my arm.<sup>5</sup>

Although we cannot perform a typical ambiguity test for the word ‘pain’ here, it is plausible that ‘pain’ is a word with two separate but related meanings, marked by distinct syntactic properties. When the word follows the preposition ‘in’, it indicates a state of a *person*. When it is modified by a prepositional phrase such as ‘in the arm’, it indicates a state of a *body part*. Indeed, given the polysemy displayed by the other three primary pain terms, this would seem to be the best way to explain the above two uses of the word ‘pain’. Hence all of the four primary pain terms seem to be polysemous.

### 3. Resolving the paradox of pain

The polysemy of primary pain terms allows us to give a response to the alleged paradox of pain.

Concepts are constituents of thoughts, just as the meanings of words are constituents of the meanings of sentences. Ambiguous words, which fail the

5 Arguably (15) is logically equivalent to ‘My arm hurts’, and in the case of (15), the word ‘pain’ should be understood as referring to a state of the body part, not a thing in the body part (Hyman 2003, Liu 2020).

relevant linguistic tests (as we saw in the previous section), make different contributions to the meanings of the sentences that they feature in and thus express distinct concepts. The polysemous word ‘expire’ expresses at least two distinct concepts – the concept of *dying* and the concept of *becoming invalid*. Similarly, the polysemy of pain terms, which can apply to both persons and body parts, reflects two separate but related folk concepts of pain. One concept tracks the so-called ‘mental conception of pain’, which reflects pain as a mental state of a person or a sentient being. The other concept of pain tracks the ‘bodily conception of pain’, which treats pain as a state of a body part.<sup>6</sup>

Once we see that there are two concepts of pain rather than one, the air of paradox dissolves. To motivate the paradox of pain, we would have to assume that there is a single concept of pain. But, as we saw, we should not think that there is only a single folk concept of pain. So there is no paradox. The intuition that there is one results from a failure to distinguish the two concepts of pain.

However, my proposal faces an objection. If there are two folk concepts of pain rather than one, why has this fact escaped theorists’ notice? After all, it is fairly obvious that there are at least two concepts of expiring and no one has confused the two. In response, two points can be made.

First, there is a significant overlapping of the contexts where the two concepts of pain can be appropriately tokened, which might make it hard to keep the two concepts apart. Consider again the paradigmatically polysemous word ‘expire’, which latches onto two separate but related concepts –  $\text{EXPIRE}_1$  (*die*) and  $\text{EXPIRE}_2$  (*become invalid*). The tokening of  $\text{EXPIRE}_1$  in the context, say, of John dying, does not license the tokening of  $\text{EXPIRE}_2$ . The fact that John expired<sub>1</sub> has nothing to do with whether any of his possessions, such as his student card, have expired<sub>2</sub>. Conversely, the tokening of  $\text{EXPIRE}_2$  does not license the tokening of  $\text{EXPIRE}_1$ . The fact that John’s card expired<sub>2</sub> by no means implies that John has expired<sub>1</sub>. Since the contexts in which the two concepts get tokened are generally distinct, this makes it easy to track the two distinct concepts of expiring.

In contrast, pain concepts seem to be different. On the current proposal, there are two concepts of pain –  $\text{PAIN}_m$  (*the mental concept*) and  $\text{PAIN}_b$  (*the bodily concept*). The tokening of  $\text{PAIN}_b$  generally licenses the tokening of  $\text{PAIN}_m$ . The fact that Sally has a  $\text{pain}_b$  in her leg typically means that she is undergoing an experience of  $\text{pain}_m$ , that is, she is in  $\text{pain}_m$ . Since pain experiences are often associated with specific body parts, the tokening of  $\text{PAIN}_m$  often licenses the tokening of  $\text{PAIN}_b$ . The fact that Sally is in  $\text{pain}_m$  often means that a part of her body has  $\text{pain}_b$ .<sup>7</sup> Such an overlap makes it hard to keep apart the two concepts

6 It is possible that one concept of pain derives from the other, just as one meaning of a polysemous word can derive from another (e.g. for ‘expire’ – the meaning ‘become invalid’ derived from the meaning ‘die’). But the direction of this derivation requires further investigation.

7 Phantom pain would constitute a case where one has  $\text{pain}_m$  but not  $\text{pain}_b$ .

of pain. So it is no surprise that theorists have overlooked the solution for the paradox of pain put forward here.

Second, polysemy can come in degrees, and the multiple concepts that polysemous words express might be regarded as distinct to different extents (Deane 1988, Tuggy 1993). Consider the verb ‘paint’, which could mean ‘cover a surface with paint’ or ‘depict or produce a picture’. Arguably, the two concepts expressed by the verb ‘paint’ are not so distinct – both involve the idea of applying coloured substance onto a surface (see Tuggy 1993). Consider also the type of polysemy known as *pros hen* (i.e. focal meaning) (Aristotle 1985: Γ2), where the related meanings of a polysemous word stem from a core meaning (Sennet 2016). For instance, the word ‘healthy’ has a core meaning which applies to things that are capable of having health, as in ‘John is a healthy person’. We also have expressions like ‘healthy blood’ (something a healthy person produces) and ‘healthy food’ (something that is conducive to health), where these related meanings of the word derive from the core meaning. One might easily mistake cases of *pros hen* for expressing unified concepts when they should be taken to express related but distinct concepts. So, insofar as the concepts expressed by polysemes are disparate to different extents, the two concepts expressed by the word ‘expire’ might be regarded as more disparate than those reflected by primary pain terms. Again, therefore, it is not surprising that the polysemy of the latter has been overlooked.

#### 4. Conclusion

As we have seen, primary pain terms – that is, ‘pain’, ‘sore’, ‘ache’ and ‘hurt’ – are arguably polysemous, applying to both a person and a person’s body part. The result of this linguistic analysis has direct implications for the alleged paradox of pain. *Pace* some theorists (e.g. Hill (2005, 2017)), there is no reason to think that there is a genuine paradox of pain. This is not because people only have a bodily conception of pain (*pace* Reuter et al. 2014, Reuter and Sytsma 2020), nor because there is a single folk concept of pain which is ‘polyeidic’ (*pace* Borg et al. 2020).<sup>8</sup> It is because there are at least two concepts of pain – mental and bodily. Once we recognize this fact, the puzzle dissolves.

Relatedly, there has been an important debate in the philosophy of pain between those who think pains are bodily conditions (see e.g. Hyman 2003, Massin 2017, Reuter et al. 2019) and those who think that pains are feelings (see e.g. Kripke 1980: 154, Lewis 1980, Tye 2017). But, given the discussion in this paper, the dichotomy seems to be ill-posed. I have argued that there are two separate but related concepts of pain manifested in our ordinary talk. So,

8 The current proposal is substantively distinct from the polyeidic view, whose proponents explicitly assume that there is a single shared concept of pain.



in asking what the nature of pain is, we ought first to distinguish between two notions of pain.<sup>9</sup>

### *Funding*

This work was supported by the Oxford Philosophical Fellowship Fund Postdoctoral Award.

*University of Hertfordshire*  
*michelle.yb.liu@gmail.com*

### *References*

- Aristotle. 1985. *Metaphysics*. In *The Complete Works of Aristotle: The Revised Oxford Translation* vol. 2. ed. J. Barnes. Trans. W.D. Ross. 4.2,1003a33-37. Princeton, NJ: Princeton University Press.
- Aydede, M. 2013. Pain. In *The Stanford Encyclopedia of Philosophy*, ed. E. N. Zalta. <http://plato.stanford.edu/archives/spr2013/entries/pain/>, last accessed 20 August 2020.
- Borg, E., R. Harrison, J. Stazicker and T. Salomons. 2020. Is the folk concept of pain polyeidic? *Mind and Language* 35: 29–47.
- Chomsky, N. 1957. *Syntactic Structures*. The Hague: Mouton.
- Deane, P. 1988. Polysemy and cognition. *Lingua* 75: 325–61.
- Fabrega, H., Jr, and S., Tyma. 1976. Language and cultural influences in the description of pain. *British Journal of Medical Psychology* 49: 349–71.
- Hardcastle, V.G. 2015. The perception of pain. In *The Oxford Handbook of the Philosophy of Perception*, ed. M. Matthen, 530–41. Oxford: Oxford University Press.
- Hill, C. 2005. Ow! The paradox of pain. In *Pain: New Essays on the Nature of Pain and the Methodology of Its Study*, ed. M. Aydede, 75–98. Cambridge, MA: MIT Press.
- Hill, C. 2017. Fault lines in familiar concepts of pain. In *The Routledge Handbook of Philosophy of Pain*, ed. J. Corns, 60–69. New York: Routledge.
- Hyman, J. 2003. Pains and places. *Philosophy* 78: 5–24.
- Kripke, S. 1980. *Naming and Necessity*. Cambridge, MA: Harvard University Press.
- Lakoff, G. 1970. A note on vagueness and ambiguity. *Linguistic Inquiry* 1: 357–59.
- Lewis, D. 1980. Mad pain and Martian pain. In *Readings in the Philosophy of Psychology*, ed. N. Block, 216–22. Cambridge, MA: Harvard University Press.
- Liu, M. 2020. The intuitive invalidity of the pain-in-mouth argument. *Analysis* 80: 463–74.
- Massin, O. 2017. Bad by nature: an axiological theory of pain. In *The Routledge Handbook of Philosophy of Pain*, ed. J. Corns, 321–33. New York: Routledge.

9 I would like to thank two anonymous reviewers for their valuable comments. Thanks also to the audience at the European Online XPhi Conference 2020 where this paper was presented. I am also indebted to the following people for helpful feedback and discussions: Emma Borg, Sabrina Coninx, Martin Davies, Domi Dessaix, Karl Egerton, Sarah Fisher, Jessica Keiser, Luke King-Salter, Dominik Kobos, Guido Löhr and Kevin Reuter.

- Reuter, K., D. Phillips and J. Sytsma. 2014. Hallucinating pain. In *Advances in Experimental Philosophy of Mind*, ed. J. Sytsma, 75–99. London: Bloomsbury Press.
- Reuter, K. 2017. The developmental challenge of the paradox of pain. *Erkenntnis* 82: 265–83.
- Reuter, K., M. Sienhold and J. Sytsma. 2019. Putting pain in its proper place. *Analysis* 79: 72–82.
- Reuter, K. and J. Sytsma. 2020. Unfelt pain. *Synthese* 197: 1777–801.
- Reznikova, T., E. Rakhilina and A. Bonch-Osmolovskaya. 2012. Towards a typology of pain predicates. *Linguistics* 50: 421–65.
- Sennet, A. 2016. Ambiguity. In *The Stanford Encyclopedia of Philosophy*, ed. E. N. Zalta. <https://plato.stanford.edu/archives/spr2016/entries/ambiguity/>, last accessed 20 August 2020.
- Tuggy, D. 1993. Ambiguity, polysemy, and vagueness. *Cognitive Linguistics* 4: 273–90.
- Tye, M. 2017. Are pains feelings? *The Monist* 100: 478–84.

## *Is credibility a guide to possibility? A challenge for toy models in science*

YLWA SJÖLIN WIRLING 

### 1. Credible toy models

A typical scientific model is an artificial system, composed of a concrete or mathematical structure with an intended interpretation, that scientists use to learn about the world. For example, one can learn about the behaviour of an aircraft by studying a scale model in a wind tunnel; or about meteorological phenomena by manipulating a computer simulation, interpreted in terms of the weather system.

*Toy models* are very abstract, simple and highly idealized models. Stock examples include [Schelling's checkerboard \(1971\)](#), [Akerlof's 'market for lemons' \(1970\)](#), [Maynard Smith and Price's Hawk-Dove game \(1973\)](#) and the Lotka–Volterra model of predator–prey dynamics. Toy models of, for example, markets or populations, are very dissimilar to, and exclude factors known to influence, actual markets or populations, and contain idealizations not susceptible to de-idealization methods.

Some toy models are 'embedded' in an empirically confirmed theory, and their epistemic merit piggybacks on the merits of that theory. But many prominent toy models, including the examples above, are not embedded but *autonomous* ([Reutlinger et al. 2018](#)). Scientists evidently find autonomous toy models epistemically useful, despite a lack of framework theory, but their epistemic merit presents a puzzle. What, and how, can they teach us about the world?