

University of Wollongong

Research Online

Faculty of Engineering and Information
Sciences - Papers: Part A

Faculty of Engineering and Information
Sciences

1-1-2016

Information-how

Nir Fresco

University of Wollongong, nfresco@uow.edu.au

Follow this and additional works at: <https://ro.uow.edu.au/eispapers>



Part of the [Engineering Commons](#), and the [Science and Technology Studies Commons](#)

Recommended Citation

Fresco, Nir, "Information-how" (2016). *Faculty of Engineering and Information Sciences - Papers: Part A*. 5201.

<https://ro.uow.edu.au/eispapers/5201>

Research Online is the open access institutional repository for the University of Wollongong. For further information contact the UOW Library: research-pubs@uow.edu.au

Information-how

Abstract

The distinction between knowledge-how and knowledge-that has long been debated in the literature. This distinction can, arguably, be better understood in terms of a more fundamental distinction between information-how and information-that. Information-how is prescriptive and informs a cognitive agent about which action(s) can be performed to achieve a particular outcome. Information-that is descriptive and informs the agent about events, objects and states of affairs in the world. Since the latter has received more attention in the literature, this article focuses on the former. We offer a new account of information-how that answers the question: 'When do data qualify as information-how?'

Disciplines

Engineering | Science and Technology Studies

Publication Details

Fresco, N. (2016). Information-how. *Australasian Journal of Philosophy*, 94 (1), 58-78.

INFORMATION-HOW[#]

Nir Fresco

Abstract. The distinction between knowledge-how and knowledge-that has long been debated in the literature. This distinction can, arguably, be better understood in terms of a more fundamental distinction between information-how and information-that. Information-how is prescriptive and informs a cognitive agent about which action(s) can be performed to achieve a particular outcome. Information-that is descriptive and informs the agent about events, objects and states of affairs in the world. Since the latter has received more attention in the literature, this article focuses on the former. We offer a new account of information-how that answers the question: ‘When do data qualify as information-how?’.

Keywords: semantic information, knowledge-how, cognitive agent, goal, action, knowledge-that

1. Introduction

Insofar as information-processing operations underpin cognitive processes, including learning, decision-making and orientating, it is useful to understand the different types of information and the roles they play in cognition. Two important types of information at the cognitive level are information-*that* and information-*how*. The former is about *states of affairs* whilst the latter is about *action selection*. Information-*that*, such as ‘Not all birds can fly’, is used by agents to represent, rather than just react to, their environment. Information-*how*, such as ‘To activate the fire alarm, break the glass and press the button’, informs the agent about which action(s) can be performed to achieve a particular outcome.

We focus on information-how and, specifically, on the question ‘What makes something information-how?’. The analysis of information-how is warranted so as to allow a discrimination of this type of information from other types and be able to compare different instances of it. A better understanding of the information-that/information-how distinction has a potential payoff in the ongoing knowledge-that/knowledge-how debate (cf. Hartland-

[#] This is a preprint of an article whose final and definitive form is published in the *Australasian Journal of Philosophy* 2015 (<http://www.tandfonline.com/doi/abs/10.1080/00048402.2015.1022561>). The *Australasian Journal of Philosophy* is available online at: <http://www.tandf.co.uk/journals/>.

Swann 1956; Hetherington 2011; Noe 2005; Ryle 1949; Stanley and Williamson 2001). There has already been a substantive debate among scholars on the nature of information-that (cf. Dretske 1981; Fetzer 2004; Floridi 2011a; Scarantino and Piccinini 2010), yet this is not the case with information-how. Moreover, since information-how is fundamental to agent-based models used, for example, in artificial life (Silverman 2007), social sciences (Billari et al. 2006; Cederman 2001) and business processes (Ghose et al. 2013), it has a potential practical payoff for the underlying agent-based framework.

We offer a new account of information-how whilst taking *data* as the starting point for understanding information. Both information-that and information-how can, but need not, have a propositional form. A map of Sydney, for example, can convey information-that about the distance between two places in the city. Similarly, information-how can be conveyed by other non-propositional means, such as by demonstrating how to perform a particular action. Information in the data-centered sense is that which is central to *cognition*; after all, cognitive agents traffic in physical *signals* they receive from and transmit to the world. Accordingly, our main question can be reformulated as ‘When do data qualify as information-how?’. The present analysis focuses on information at the *symbolic* level, which includes symbolic constructs such as natural and programming languages, logic, maps and diagrams.

We begin (Section 2) by briefly surveying the knowledge-how/knowledge-that debate as a background to the present enquiry. In Section 3, we offer an account of information-how. Section 4 is devoted to addressing two possible objections. Section 5 critically compares the present account with two other accounts of information-how. Section 6 concludes this article with some suggestions for future work.

2. Background: Knowledge-how versus Knowledge-that

The distinction between *knowledge-how* and *knowledge-that* has long been debated in the literature. The former concerns knowledge how to do something and is manifested in the performance of a skill, such as riding a bike or playing a piano. The other concerns knowledge of facts, for example, that ‘Novak Djokovic was the winner of the 2015 Australian Open’ and is cognitively representable. In the next three subsections, we briefly review three competing views in this debate, respectively. The fourth subsection connects the knowledge-how/knowledge-that distinction with the information-how/information-that distinction.

2.1. *Knowledge-that and knowledge-how as distinct types of knowledge*

A well-known proponent of the view that knowledge-that and knowledge-how are distinct types of knowledge (DTK) was Gilbert Ryle. He described knowledge-how as follows:

“Knowing *how* [...] is a disposition, but not a single-track disposition like a reflex or a habit. Its exercises are observances of rules or canons or the application of criteria, but they are not tandem operations of theoretically avowing maxims and then putting them into practice” (Ryle 1949, p. 46, italics original). “Learning *how* [...] is not like learning *that* or acquiring information[-that]. Truths can be imparted, procedures can only be inculcated” (Ryle 1949, p. 59, italics original).

To show that knowledge-how cannot be reduced to knowledge-that, Ryle famously argued that the former cannot be dependent on the latter without entering an infinite cycle.

“The consideration of propositions is itself an operation the execution of which can be more or less intelligent [...] But if, for any operation to be intelligently executed, a prior theoretical operation had first to be performed and performed intelligently, it would be a logical impossibility for anyone ever to break into the circle” (Ryle 1949, p. 30).

The intelligent execution of an action need not be either preceded or constituted by consideration of some information-that p , let alone by knowledge of p 's being true. One dominant reaction to Ryle's analysis was to accept the DTK view. For “[k]nowledge-that is a cognitive state in which one accurately represents [...] some aspect of reality [...] Knowledge-how is an ability, a non-necessarily-cognitive capacity to do or act” (Hetherington 2011, p. 27).

2.2. *Knowledge-that as a subtype of knowledge-how*

According to the second view, all knowledge is knowledge-how (KH). According to John Hartland-Swann, knowledge-that in ordinary discourse, such as in “I know that the French for ‘boy’ is ‘garçon’” is reducible to knowledge-how by regarding the verb ‘to know’ as dispositional (1956, p. 112). On his view, the only way to maintain a robust distinction between these two types of knowledge is “to abandon the idea that ‘knowing’ is a purely dispositional concept” (Hartland-Swann 1956, p. 114). However, since the price for such a move is too high, the only viable alternative is “to deny that knowing *that* is basically different from knowing *how*” (ibid). Knowledge-that p is reducible to knowledge-how, since the former is simply the *capacity* to represent or state correctly how it is that p is known.

A variant of the KH view is offered by Luciano Floridi in his defence of constructionist epistemology. On his view, knowledge-how is more fundamental, in terms of data input, than knowledge-that. “[U]ser's and maker's knowledge are complementary types of know-how, which must be joined together in order to reach full and useful *episteme*” (Floridi 2011b, p.

288). Rather than knowing a phenomenon or an object by merely passively observing it and then forming a cognitive representation of it as a “copy”, knowledge is gained, according to Floridi, by being able to *produce* it and *reproduce* it.

A recent defence of the KH view has been offered by Stephen Hetherington arguing that all knowledge-that is essentially reducible to knowledge-how. At the core of his theory (‘Practicalism’) is the *knowledge-as-ability* hypothesis that can be described as follows. “Your knowing that *p* is your having the ability to manifest various accurate representations of *p*. The knowledge as such is the ability as such” (Hetherington 2011, p. 42, italics original). Hetherington argues that it is epistemologically arbitrary to insist that knowledge-that is a belief, and not another cognitive member of the knower’s “epistemic diaspora”, such as sensing that *p*, reasoning that *p* and asserting that *p* (2011, pp. 36–37). A consequence of Practicalism is that if knowledge is indeed an ability, knowledge comes in degrees, since abilities come in degrees.

2.3. *Knowledge-how as a subtype of knowledge-that*

According to the third view, all knowledge is knowledge-that (KT). One well-known argument in support of this view relies on the idea that ascriptions of knowledge-how are ascriptions of propositional knowledge, in a manner akin to ascriptions of knowledge-*who*, knowledge-*what*, knowledge-*when* and knowledge-*where* (Stanley and Williamson 2001, p. 420). In essence, knowledge-how consists in knowing the answer to a question that is syntactically embedded in the sentence that attributes knowledge-how to the subject. So, to say, for example, that Bob knows how to ride a bike is akin to saying that he knows *that* such-and-such is a way for him to ride a bike. “Know-how is a positive doxastic attitude [...] toward a method in which [...] movements [or actions] are represented in a way that the agent can execute at will in a broad range of circumstances” (Perry 2003, p. 154).

2.4. From knowledge-how/knowledge-that to information-how/information-that

At least *prima facie*, the acquisition of knowledge-how seems to depend ultimately on the agent being informed about how to do something. (This idea resonates, for example, in (Poston 2009) where the “intelligence base” for skilfully performing an action is identified with the set of instructions one acquires on the basis of her training.) An agent can, but need not necessarily, be informed about how to *F* by another agent or some external information source. After all, she may learn how to *F* by other means – predominantly by trial and error – without a teacher imparting any information-how to *F*.

Still, consider the case of a teacher, Bob, imparting information-how to ride a bike to Alice. Suppose that he tells her ‘To ride a bike, do it like this’ followed by a live riding demonstration. Bob’s demonstration seems to convey information-how to ride a bike without necessarily conveying information-that. Any information-that verbally conveyed might even be false, whilst the demonstration is performed correctly.¹ This invites the question whether knowledge-how obtained from information-how is also susceptible to Gettierisation (for a discussion see (Poston 2009) in reply to (Stanley and Williamson 2001, p. 435)). In knowledge-that, Gettierisation arises from a disconnect between the veridicality of information-that *p* and the corresponding justified belief-that *p*. In knowledge-how, the disconnect is between the information-how and its success in bringing about the target result. Yet, in the latter case, both the effectiveness and the usability of information-how – that are discussed in Sections 3.3 and 3.4 – should be considered for a complete analysis of potential Gettierisation of knowledge-how.

Additionally, although Bob is manifesting knowledge-how to ride a bike, it is information-how being transmitted. Information-that is intimately linked to the agent’s belief-that about

¹ An analogue can be found in (Stanley and Williamson 2001, p. 435) and (Poston 2009) where (another) Bob wishes to learn how to fly in a flight simulator and is trained by an imposter instructor, who intends to give Bob wrong advice.

some state of affairs², whereas information-how is linked to her skill(s) and action selection. By demonstrating how to ride a bike, what is transmitted is information-how about appropriate action selection under the given circumstances for two main reasons. (1) The motor *skills* to successfully ride the bike, which are critical for knowledge-how, are not easily transferrable. Skills, typically, develop *gradually* and even having a skill fully developed does not entail the *ability* to correctly *describe* it (we return to this point in Section 4.2). (2) Understanding and recognition of Bob observed actions afford Alice the repertoire of action selection in the *actual* scenario. However, his knowledge-how to ride the bike under *unactualised* circumstances is not demonstrated.

The information-how conveyed to Alice can, but need not, give rise to her knowledge-how to ride a bike. Alice may try to exercise the information-how conveyed, but may not yet *succeed* in riding a bike. For her riding skills have not fully developed thus far (and she does not have the requisite transferrable skills, such as riding a horse). She may not be able to successfully exercise some of the instructions to make the required moves correctly. Her learning how to ride a bike is accomplished by either observational learning (as the discovery of affordances for particular goal-based actions and the transfer of actions on affordances) or by repeated trial and error.³

Whether or not the information-that/information-how distinction is fundamental to the knowledge-that/knowledge-how debate warrants further discussion (cf. Dretske 1981; Floridi 2011a, Chapter 12; Hetherington 2011, p. 71). We say more about these two parallel distinctions in dealing with possible objections in Section 4. If the view suggested in this article is correct, information-how is prior to knowledge-how (at least in many paradigmatic

² Information-that should be distinguished from the content of belief-that, and as such the former is not superfluous. One may receive the information *p*, but choose to ignore it, due to former entrenched beliefs (perhaps, with good evidence) that $\neg p$. Moreover, one cannot simply *decide* to believe that *p*, even when she is informed that *p*.

³ See (Hartley et al. 2008) for a discussion of the two learning methods and some reported experimental results.

cases). A complete theory showing when information-how upgrades to knowledge-how is undoubtedly needed. But first we offer an account of information-how.

3. What is Information-how?

There has been less discussion about the information-that/information-how distinction than about the knowledge-that/knowledge-how distinction. On some functional-evolutionary definitions of semantic information, the former distinction becomes very blurry. That is the case, for example, when ‘functional’ means that signals received by either a human- or natural-selection designed system play a causal role that “usually contributes to the goal-oriented behavior of this system” (Jablonka 2002, p. 582). An apple pie recipe and a piece of software are instances of functional-evolutionary information for a cook and a computer, respectively, in a manner akin to the appearance of black cloudy sky leading to the shelter-seeking action of an observing ape (Jablonka 2002, p. 580). Nevertheless, insofar as we seek to understand the role information processing plays in cognitive tasks on an *ontogenetic* (lifetime of the individual) scale, rather than on a *phylogenetic* (evolutionary) scale alone, the information-how/information-that distinction seems worth preserving.⁴

3.1. Data, semantic content and truth

Since our focus is on information processing by cognitive agents, we take *data* to be the starting point for understanding information. Our main question is ‘When do data qualify as information-how?’. To answer it, we need a working definition of ‘data’ and would be better served by adopting a relatively broad definition.

For simplicity, data are broadly defined here as *physical discontinuities in the world* in the spirit of the Floridian definition (a datum is a variable x that is distinct from another variable

⁴ Similarly, Floridi notes that in some contexts, “there is the danger that the concept of biological information might lose its [...] concrete *procedural* sense, and silently acquire an increasing *semantic* sense” (2010, p. 81, italics added).

y). According to Floridi, data can be treated as “diaphora *de re*” – the lack of uniformity in the real world, “diaphora *de signo*” – the lack of uniformity between the *perception* of at least two signals, or “diaphora *de dicto*” – the lack of uniformity between at least two symbols (Floridi 2011a, pp. 85–86). Thus defined, nothing is strictly *a datum* without its counterpart. The differences between high and low voltage and the letters α and β in the Greek alphabet are examples of data.⁵

Moreover, according to Floridi, for data to qualify as *semantic content* they need to have two additional properties: (1) be well formed, and (2) be meaningful (Floridi 2010, p. 20). *Well-formedness* means that the “data are rightly put together, according to the rules (syntax) that govern the chosen system, code or language being used. Syntax here must be understood broadly, not just linguistically” (Floridi 2010, pp. 20–21). *Meaningfulness* means that the “data must comply with the meanings (semantics) of the chosen system, code or language in question. [... They] can be meaningful independently of an informee [and need not be] necessarily linguistic” (Floridi 2010, pp. 21–22). For present purposes, we adopt Floridi’s definition of semantic content as *meaningful*⁶ *structured data* (2011a, p. 84). Although there are problems with this definition, such as ‘how are the data structured?’ and ‘how can data be meaningful absent a consumer of information?’, they exceed the scope of this article.

The distinction between information-that and information-how is analogous to the distinction between declaratives and imperatives.⁷ Information-that is truth-evaluable and is used by agents to *represent* and *form beliefs* about, rather than to react *externally* to, their environment. Information-how is, arguably, truth-neutral. Yet, the property of *being truth-*

⁵ There are problems with Floridi’s definition, such as whether data persist or what qualifies as “distinct variables”, but they are not discussed here. A plausible alternative is defining a datum (or a cue) as any act or structure that can be received by the system concerned (Fresco et al., unpublished).

⁶ We take meaningfulness to be determined by the behavioural effect of the data on the receiver.

⁷ The information-how/information-that distinction, arguably, also applies at a non-symbolic level in some non-human animals, such as apes, parrots and corvids. In relatively simple organisms this distinction is harder to discern (cf. Millikan 2004, pp. 17–18, for the case of distance messages about nectar sources in dances of honeybees). This idea is further developed in Fresco et al. (unpublished) in the context of dealing with the relation between information and (general) selection.

evaluabile depends on one's choice of a theory of truth, such as a correspondence theory, a coherence theory or a pragmatic theory. (Truth is typically taken as correspondence to reality.)

What is our background theory of truth, then? Floridi proposes the *correctness theory* of truth (CTT) for fixing the veridicality of semantic content giving rise to information-that. Simplistically, on CTT, the semantic content *i* an agent holds can be polarised into a query, *Q*, and a result, *R*, “qualified by a specific context, a level of abstraction [(LoA)] and a purpose” (Floridi 2011a, p. 182). The agent has a model *m* of a system *s* that is produced by *i* relative to *Q* and the relevant LoA⁸. The data underlying *i* are true *iff* *m* can be used by the agent as a proxy (a term borrowed from computer science) for *s*. The agent's access to *s* is mediated by *m*. The alethic value of *R* can, but need not, be that of correspondence to reality (Floridi argues for a non-correspondentist sense of correctness).

What matters for present purposes is a common ground between CTT and the pragmatic approach to truth adopted for our analysis. CTT has a pragmatic sense (not the pragmatic-theory-of-truth sense): the *agent* is essential to determining the truth of *i*. The pragmatic approach to truth herein is guided by the agent's successful negotiation with her environment and her ability to make accurate predictions about the future relative to a given LoA and a particular purpose. Undoubtedly, the correspondence of information to the external world is conducive to the practical success of the agent's negotiation with her environment. But correspondence to reality as a *general, all-purpose* criterion of truth is problematic. (Is there only *one* way to correspond? Besides, some propositions do not have easy-to-verify correspondence relations.) Similarly, on CTT, the better *m* can be used to successfully predict

⁸ An LoA is an object of study (e.g., a system) consisting of a finite, non-empty set of observables, which are *interpreted* typed variables (Floridi 2011a, p. 52). A typed variable is a variable qualified to hold only some declared kind of data. In this context, ‘interpreted’ means that the typed variable represents some feature of the system under consideration.

some aspect of s , the smaller the gap is between its forecast and what happens in s (Floridi 2011a, p. 202).

3.2. *An outline of information-how*

Whatever additional features semantic content needs to have to qualify as information-how they should yield truth-neutral content. Nevertheless, information-how *is truth seeking* in the sense that the action(s) prescribed should be successful as is discussed below. These features are reflected by the following principle:

The Principle of Information-How (PIH): semantic content is an instance of information-how *iff* its fulfilment through action *can* systematically produce a particular outcome in a given context.⁹

First, we note two basic desiderata that are suggested by PIH: the performability and lack of ambiguity of the underlying semantic content. Consider the instruction ‘Pick up the bat and Do not pick up the bat!’ and suppose that on the mat in question there are exactly two objects: an injured *flying bat* and a *cricket bat*. This instruction *cannot* be followed. Unless the two occurrences of ‘bat’ in that instruction have *different* referents (the flying bat and the cricket bat), there is no possible action the performance of which can *effectively* yield the specified outcome. Similarly, if an instruction cannot be followed in principle, for example, ‘Open the window and Do not open the window!’, there is little point in instructing an agent to do it.¹⁰ There is no right condition under which *any* agent can follow a contradictory

⁹ We note the similarity to the sufficient condition articulated in the instructional information processing account of computation (Fresco and Wolf 2014). The present definition offers a more precise characterisation of information-how in the context of cognition.

¹⁰ Interestingly, according to deontic logic, such a contradictory action, which is never *performed* by an agent, is, nevertheless, not necessarily forbidden (von Wright 1951, p. 11).

instruction. That instruction still qualifies as information though, much like a logical contradiction qualifies as (false) semantic information.

Next, consider the instruction ‘Pick up the bat on the mat!’, PUB. The content of PUB is *ambiguous*. Whichever of the two corresponding actions the recipient performs, PUB in isolation cannot be used to determine whether the executed instruction was followed as prescribed. Nevertheless, an agent following *PUB* can repeatedly lead to the same outcome (e.g., picking up the cricket bat). But when given to *different* agents it is less likely to be repeated by *all* agents to produce the *same* outcome. Ambiguity in the prescription of action hinders systematicity. It is the *prescription* of the action that has to be unambiguous, rather than the *action* itself.

Whilst PIH suggests that the prescription of the action(s) be unambiguous, the prescription of the action(s) *can be* vague. Consider a lieutenant commanding one of her soldiers to inform the cook to have lunch ready at noon for the platoon coming back from active duty. Whether the soldier informs the cook directly or via some proxy, by calling the cook, texting him or telling him in person is unspecified. The means of informing the cook is vague, and the command can, nevertheless, be effective. This should be unsurprising given that a requirement for excessive precision in specifying an action may lead to infinite regress. Background assumptions are made in exercising information-how (just as in processing information-that).

When a sender, *S*, of information-how, *x*, instructs an agent, *R*, to do *y* and *R* does so, *R* has acted (i.e., satisfied $x \rightarrow \text{action}$) as a result of being instructed by *S* to do *y*. As such, information-how has two types of relations to the set of actions it encompasses. The first is a *logical* relation: it specifies and enjoins the set of actions *y* as prescribed by *x* (Hamblin 1987, p. 93) to be performed by an agent. *x* connects *R* and the object(s) to be acted on. It is a

logical relation insofar as information-how is understood as a relation between an agent and the action(s) prescribed.

The second relation is causal: x acts as a factor in getting the set of actions realised by R (ibid). It plays a causal role in affecting the agent's actions insofar as in the absence of x R would not have performed the action(s) y as prescribed by x (at least not systematically). This causal relation also presupposes that R not only wishes to achieve the underlying goal by satisfying x , but also has the *capacity* to do y , in the sense of R possessing the physical power required to do y . R not 'having the capacity' means that the "conditions whose satisfaction is necessary for [R] to exercise [R 's] ability are not satisfied" (Noe 2005, p. 283). The loss of legs, suffering an attack of vertigo and lack of access to a bike – each renders R incapable of riding a bike.

Information-how is in essence imperative. Unlike information-that, it prescribes a set of actions to be performed by the informee. As such, it cannot be straightforwardly evaluated alethically, but rather as *effective* or *ineffective* in achieving a particular outcome. Imperatives are evaluated relative to their utility in achieving some outcome or goal. The conflation of information-how 'being-satisfied' and 'being-true' is simply the result of viewing the satisfaction of information-how as singling out the set of possible worlds in which the relevant state of affairs obtains (Hamblin 1987, p. 152). The performance or non-performance of an action prescribed by information-how is comparable to the truth-value of information-that. (Consider again a contradictory instruction.) But neither an instruction nor the action information-how prescribes is either true or false *in itself*.¹¹

To clarify the relation between information-how and imperatives consider three similar messages sent to some receiver.

¹¹ A similar point was stressed in (von Wright 1951, p. 15): unlike the alethic modality, the deontic modality has no logical connection with matters of fact. For nothing follows from the performance (or non-performance) of an action regarding the action's permitted, obligatory or forbidden character.

- 1) ‘Activate the fire alarm!’.
- 2) ‘To activate the fire alarm, activate the fire alarm!’.
- 3) ‘To notify all residents of the fire, activate the fire alarm!’.

At least *prima facie*, it might seem that only the second and third messages qualify as information-how, as only they explicitly specify the particular outcome. Against this specific outcome the effectiveness of the information-how can be evaluated. Yet, a closer examination reveals a greater similarity between the first and second message. These are both evaluated against the same condition obtaining (the fire alarm being activated). The second message specifies the outcome explicitly and, is, therefore, less efficient in terms of the amount of information it encodes.

The upshot is that only the third message is an instance of *effective* information-how. Relative to “the fire alarm having been activated” LoA – the first two messages convey *ineffective* information-how. For neither message specifies *how* to activate the fire alarm. The third message, too, does not specify how to activate the fire alarm. But, relative to another LoA: “informing interested parties of the fire that broke out”, it does specify a relevant action. A recipient obeying an instruction does not yet amount to her acquiring effective information-how without the extra context of what goal/outcome should be accomplished. The outcome can, but need not, be explicitly included in the message. Yet, it has to be known in the context in which the information-how is used/given.

PIH only requires that the prescribed action *can systematically* produce, rather than necessarily produce, a particular outcome. This weaker requirement leads to PIH being more inclusive in terms of classifying semantic content as information-how than if the *can* modality were dropped. Consider the message ‘To lower the room temperature to 25°C – whilst the air conditioner is set to 25°C – close all the windows in the room’, RT25. Suppose that the current room temperature is greater than 25°C. Roby, the robot, seeks to confirm the

relevant state of affairs (in this sense, RT25 is truth seeking in terms of satisfying some pre-conditions) by checking that the air conditioner, which is not malfunctioning, is set to 25°C. It proceeds to successfully close all three windows in the room.

As long as the state of affairs *room temperature being 25°C* is attainable in principle in the scenario described by fulfilling RT25, it qualifies as information-how. Yet, the set of actions prescribed by information-how should be such that it can systematically, rather than fortuitously, produce the particular outcome. RT25 implies that the initial room temperature is greater than 25°C. But if the room temperature drops to 25°C because of a drop in the *external* temperature, *closing* the windows may be detrimental to lowering the room's temperature. They should be opened to allow air circulation to cool the room.

Such example is commonplace in the design and implementation of artificial agents and robots. Agents use information-that about states of the environment and information-how for procedural planning. A triggering event leads to the invocation of context-sensitive plans. The execution of the plan produces a series of actions that affect the agent's environment. Deontic logic (and, more broadly, dynamic logic) is used as a means of reasoning about the actions of artificial systems, including the possession, maintenance and updates of goals and commitments (Meyer 2000). The specification of the effect that the room temperature is greater than 25°C indicates that Roby should see to it that the pre-condition of the plan with respect to that effect is realised. Deontic logic offers the formal means of negotiating (conflicting) obligations, prohibitions and permissions (von Wright 1951) (e.g., when a non-violation of obligation x is preferred over the satisfaction of obligation y).

The action-directed nature of information-how discriminates semantic content that qualifies as information-how from information-that. And once an instance of semantic content has been identified as information-how, it can be further qualified as good/effective

and/or usable. These two additional features are examined in the next two subsections, respectively.

3.3. *The quality of information-how*

The second feature of information-how is its quality, to which we have alluded above as *effectiveness*. This feature is comparable to the *veridicality* of information-that.¹² True information-that is epistemically preferable to false information-that and good/effective information-how is pragmatically preferable to poor/ineffective information-how.

When is RT25 *good* information-how? It will not suffice to simply determine whether all three windows in the room have been successfully closed (by Roby). Roby might have closed all the windows successfully, and yet the room's temperature is still greater than 25°C. Two crucial criteria for determining the quality of RT25 as information-how are (1) whether the relevant outcome (lowering the room's temperature) relative to an LoA (e.g., one that only includes *external-facing* windows) is specified/known, and (2) whether that outcome *can be* systematically attained in the given context. We have already touched upon the first criterion above. This criterion can be derived from the logical relation that information-how has to the set of actions it encompasses. RT25 explicitly specifies a particular outcome that the prescribed set of actions is to achieve. So, it is a good candidate of effective information-how. However, the particular outcome can just as well be known in context without being explicitly encoded in the information-how. (Consider the message "Activate the fire alarm!" given in a fire drill.)

The second criterion concerns the causal relation that information-how has to the set of actions it encompasses. Whilst information-how need not be *actually* fulfilled in order to qualify as good/effective, the actions prescribed need to be *satisfiable* in the given context. It

¹² Analysing the (dis)similarities in detail relative to the chosen theory of truth exceeds the scope of this article. However, a future analysis of the (dis)similarities may yield insights about the epistemic roles of information-that and information-how.

can be evaluated analytically without being realised by an agent. Consider another example from computer science: an algorithm for sorting an array of numbers. The algorithm need not be executed on a computer to be evaluated as good or poor at sorting an array. It can be evaluated analytically as performing better or worse than another algorithm relative to this task. (In complexity theory computational problems are classified into hardness classes and algorithms for solving these problems are evaluated analytically for their qualitative performance.)

Possible worlds semantics can be used as a comparative metric for assessing the quality of information-how. The more possible worlds there are where RT25 can lead to the room temperature being 25°C (or cooler), the better RT25 is with regard to the context in question. If it is not physically possible for the room's temperature to drop to 25°C (i.e., there is no possible world in which that state of affairs can hold), then RT25 is not good information-how. In each possible world, factors, such as the heat generated by the number of people in the room, the power of the air conditioner relative to the size of the room and the external temperature, contribute to the prescribed actions leading to the specified result.

This shows that the quality of information-how is context-relative. The number of possible worlds under consideration varies according to context and the LoA chosen. The quality ordering of RT25 depends on its result in possible worlds closer to the actual world. Its effectiveness lessens the more distant the possible world is from the actual world. This characterisation of effectiveness makes clear comparisons possible in some cases, but not in all. Whilst a quantitative characterisation of effectiveness is worth exploring, this comparative metric provides a good starting point for assessing different instances of information-how.

3.4. *The usability of information-how*

The quality of information-how should be distinguished from – the third feature – its usability. The *usability* of information-how is comparable to the *usefulness* of information-that. Despite the truth of a tautology, it is not very useful information-that, and, on the other extreme, a contradictory instruction leading to inaction is unusable.

The question concerning the quality of RT25 in the scenario described above presupposes that the informed agent has the *capacity* to close the windows. Suppose that RT25 was issued to Shorty the robot. Shorty cannot reach the windows in order to close them: it is too short. If there were no available means in the room by which Shorty would have the capacity to close the windows (e.g., by climbing a chair), RT25 would be *unusable* as information-how for Shorty. Nevertheless, RT25 can still be *good* information-how. Also, the very same information would be *usable* for Roby.

This again shows the context- and LoA-relativity of information-how. A simple baking recipe may be completely intelligible to a six year old, but absent the capacity to use the oven, he cannot use this information-how. The baking recipe is, therefore, unusable information-how for the six year old.

3.5. *Actions and negative actions*

In analysing cognitive agents, it is important to distinguish between *actions* executed by agents and *happenings* that just happen without requiring any agent intervention. Charles Hamblin called the former *deeds* (1987, pp. 140–144). Actions, typically, have a purpose. Some actions may be deliberate and others may be accidental, in which case, it raises the question whether they should be classified as actions to start with.

To clarify this distinction, consider another example. Suppose that when Roby receives RT25 it closes two windows in the room. When it proceeds to close the last one that window already happens to be closed. Closing the first two windows can be attributed to Roby as

deliberate, purposeful actions. The closing of the third window is either a happening (e.g., the window was closed shut by a gust of wind) or an action executed by another agent (e.g., some person who was disturbed by the gusty wind). In what sense, then, would it be correct to claim that Roby *successfully* followed RT25? On one extreme, we may view RT25 as not having been executed effectively, since the action prescribed by RT25 was only partially satisfied. On the other extreme, we may view RT25 as having been executed effectively by Roby performing a *negative* action.

For PIH to hold true in practice, a *negative* action in some circumstances should be permitted as *valid*. That is, there exists the possibility of a null-action (Hamblin 1987, p. 141). When Roby is faced with the third window already being closed, it recognises that the window is in the right “closed-state”. Rather than taking a redundant action to *reopen* the window and then close it, thereby exerting energy unnecessarily, Roby performs a negative action. It leaves the state of the window unchanged thereby fulfilling RT25. The action RT25 prescribes could be reformulated as ‘close all the *open* windows in the room’. However, this latter formulation either makes some information *explicit* or adds *new* information to the original message.

A negative action should be classified as a valid action, then, when the *same outcome/goal* can be similarly achieved by minimising the effort exerted in taking no direct action. This principle is equally true for robots, biological organisms and humans alike. It accords with the general “law of least effort” that applies to cognitive as well as physical exertion. This law asserts that “if there are several ways of achieving the same goal, people will eventually gravitate to the least demanding course of action” (Kahneman 2011, p. 35). This consideration shows that a simple examination of the intended result of some information-how process may be insufficient for determining the effectiveness of information-how.

4. Two Objections and Replies

In this section, we address two objections in the next two subsections, respectively.

4.1. First objection: information-how is not a distinct type

The first objection dismisses information-how as a distinct type of information in a manner analogous to the KT view discussed in Section 2.3. Whilst instructions are not propositions and are, thus, neither true nor false, every instruction can be *translated* to a corresponding conditional. If the instruction is followed, the specific result is achieved. For example, ‘Turn the light dimmer 180° clockwise to switch off the light’ can be translated to the truth-evaluable proposition, ‘If the light dimmer is turned 180° clockwise, the light is switched off’.

Four observations are in order in reply to this objection. Firstly, as we know from Information Theory, information can be measured. It is not clear that the translation of information-how into a proposition preserves *all* the original information. Brian Skyrms, for one, analyses information *about the state of nature* (or information-that) and information *about the act* (or information-how) in the context of signalling games¹³ and learning dynamics. He argues that the definitions of information-how “are entirely analogous to those of information about the state [of nature]” (Skyrms 2010, p. 39). Skyrms adds that “probabilities of the states, probabilities of sender’s strategies and probabilities of receiver’s strategies give us unconditional probabilities of the act” (ibid). Nevertheless, as he admits, the *quantities* of these two types of information in the message *need not be the same*. Suppose, for example, that the sender selects a different signal for each state, but the recipient does not pay attention and always acts the same. In that case, there is much

¹³ In signalling games, information facilitates social coordination between a sender observing a state of the world and a receiver responding to the sender’s signals, which track the environment and motivate behaviour.

information about the states in the message but zero information about the actions (Skyrms 2010, pp. 39–40).

Secondly, the translation proposed above changes the structure of the information-how: the order of the underlying data and the relationships among them in the corresponding proposition convey different meanings. The subdivision of information-how into smaller units may result in them no longer qualifying as information-how in isolation. A recipe is a quintessential example of information-how. One can extract the instructions for baking a cake from a recipe and construct the conditional ‘If one follows instructions 1,2,3...n, one gets a mud chocolate cake’. Suppose that the second instruction is ‘beat five eggs and add two spoons of sugar’. The propositional translation of that individual instruction, ‘If one beats five eggs and adds two spoons of sugar, one is $\frac{2}{n}$ th of the way through completing the mud chocolate’, seems problematic. Unless the sub-result of each individual instruction was contained within its propositional counterpart, it would seem that some information is lost in the process. In some cases, a sub-result can straightforwardly be specified, such as ‘to prepare the sugar egg mixture for the cake’. But even if sub-results were to be included in the propositional counterparts, the *translatability* of information-how to information-that would not imply the *reducibility* of the former to the latter. For additional information has to be added in the translation process.

Thirdly, as in Bob’s bike riding demonstration, the information-how need not convey everything about how to ride a bike in terms of information-that and any information-that used may even convey something *false*. Whilst it is possible for Bob to translate every instruction or act of demonstration into a propositional counterpart, this would seem very unnatural.

Arguably, such a description of information-how lacks the requisite “imperative force” required for Alice to learn effectively from Bob. This point can also be seen more narrowly

in the distinction between the indicative and imperative moods. For example, the imperative function of ‘Open the door’ is lost in its indicative counterpart formulation ‘The door is open’. Conversely, the former may lose its indicative function, because the door is closed and may never be opened (Huttegger 2007, p. 426). Clearly, effective information-how requires a context for the set of actions prescribed relative to a specific LoA. (And the imperative ‘Open door’ is unconditional.)

Fourthly, when Alice receives an instruction from Bob to perform a particular task, there is an expectation that she would perform a particular action (if some necessary conditions obtain). Bob might send her the message ‘If the handle bar is held straight, the front wheel is aligned with the rear wheel’. Alice might have a good working memory. So, perhaps, she could even always remember this message when riding her bike whilst regularly examining the environment for any obstacles in her way. But unless the sub-result(s) of that particular instruction is (are) specified in the proposition, why should Alice be inclined to put this information to good use?

The answer is related to the ongoing debate regarding the distinction between declaratives and imperatives. An initial intuition might be that when Alice receives information-that in a propositional form (‘If the handle bar...’) she has to combine it with other indicatives to determine the most efficient behaviour (Huttegger 2007, p. 417). David Lewis distinguished among three types of signals: indicative (signal-that), imperative (signal-to) and neutral. In the indicative case, he argued, “it was up to the [sender] to *see to it that* his signal was true by choosing the correct signal to give” (Lewis 1969, p. 150, italics added). In the imperative case, “it is up to the [recipient] to make the [sender’s] signal true by responding to it correctly” (ibid). The conventional meaning of the signal-to sets an expected response that should be taken in virtue of the signal.

Information-how often leads to behaviour directly and refers to states of affairs only *indirectly*. Some state-act pair messages require the sender to *deliberate*, whereas the recipient need not deliberate, but rather *act* (Huttegger 2007, pp. 410, 417). A signal might be interpreted as information-how when the sender has to *classify* the state of affairs, whereas the recipient has to *react* (Huttegger 2007, pp. 415–416) to achieve a particular result. Imperatives, as is specified above, can often simply tell someone what to do without specifying explicitly why or how to do it. And, in cases that “the *sender* does not know what the correct action is, how could he issue a *directive*” (Zollman 2011, p. 167, italics added) or information-how?

4.2. *Second objection: information-how and belief-how*

The second objection emerges in consideration of John Perry’s classification of knowledge-how as a subtype of knowledge-that “about ‘way of’ relations [... based on] *belief-how*, which is a state that is internally like know-how, except the way-of relation doesn’t hold” (2003, p. 153).¹⁴ Bob may inform Alice about how to ride a bike, and she might even pass this information on to others without knowing how to ride a bike. For “[k]now-how is a matter of attunement to a method, not possession of a *formula describing the method*” (Perry 2003, p. 154, italics added). If the agent’s “belief-how” is supposedly true, the action should be a way of bringing about the particular result. If that is right, information-how *is* straightforwardly truth-evaluable.

However, there does not seem to be a well-founded cognitive notion of ‘belief-how’ that is akin to that of ‘belief-that’. According to Perry, the action(s) to be executed is (are) internally represented in the agent leaving “the agent attuned to the way-of relation between the executions and the [intended] results” (ibid). Such representation “can be true or false,

¹⁴ A way-of relation is a relation between accomplishments of an agent and the circumstances she is in. For Perry, ‘way-of relations’ can be described as the method by which something can be accomplished by the agent.

and is naturally regarded as a part of various concepts we have of various action” (Perry 2003, p. 155). On his view, Alice “can [be] given an excellent verbal description of the movements required for” (Perry 2003, p. 156) riding her bike. But if she cannot perform these movements, her representations are not *executable*: Alice only has a belief-how to ride a bike.

The reason for the absence of a well-founded cognitive notion of ‘belief-how’ is similar to the lack of a more general notion that subsumes the notion of ‘skill’. Skills that underlie knowledge-how are *success-related*. Knowledge-that is a kind of belief-that, which is true and satisfies some other condition(s). Beliefs-that in and of themselves do not always suffice to direct the agent’s action. Despite believing that it would rain today, Alice may not take an umbrella on her way to work. There need to be suitable dominant desires (e.g., Alice’s desire not to get wet that dominates her desire to minimise the number of items she takes to work) that together with her belief-state jointly lead to Alice acting in a certain way (e.g., taking the umbrella).

Belief-how, as it is described by Perry, cannot be comparable to belief-that in the way he suggests, since the former crucially depends on the gradual development of skills. Alice may not have the skills required to ride a bike fully developed yet. But she may still gain the “feel” of the movements required and how they contribute to maintaining her balance by observing Bob’s demonstration. If Alice is still developing that skill, she does not have the skill* to ride a bike, where skill* relates to skill as belief-that relates to knowledge-that.

Moreover, even when a skill is fully developed, many people cannot *describe* it. Being able to describe a skill is not necessary to having it (e.g., toddlers have many skills that they cannot verbalise). “While for the most part people acquire skills by imitation and trial and error, in more formal instruction we start with rules, [... which] seem to give way to more flexible responses as we [gradually] become skilled” (Dreyfus 2006, p. 46). Information-how

is “like training wheels [w]e may need [...] when learning to ride a bicycle, but [...] eventually set [...] aside” (ibid). Some cognitive scientists claim that knowledge-how is represented in a *non-conceptual* manner. Yet, “we can be aware of the goals of our actions [...] without being aware of the detail of how the [sensory]-motor system achieves [them]” (Reed et al. 2010, p. 74).

Information-how to F can be used by an agent to practise certain prescribed actions whilst she develops an understanding of the skills to F . As she becomes competent (possibly, even an expert) in performing that activity, the reliance on explicit information-how gradually subsides. “[A] master pianist who knows how to play piano [...] and] has lost her arms in a tragic accident [...] cannot play piano” (Noe 2005, p. 282), still, she can potentially convey information-how to play piano to a novice student. Even if the student possesses the physical power to exercise the information-how, she can hardly be attributed know-how to play piano yet.

Being able to exercise information-how to F is not sufficient for knowing-how to F . “The ability to give by rote the correct solutions of multiplication problems differs in certain important respects from the ability to solve them by calculating” (Ryle 1949, p. 42). (Of course, giving the correct solution for 7×6 by rote does not imply the inability of an agent to calculate it.) Mechanically exercising information-how to F does not imply an understanding of how to F . It implies that a failure to perform a *necessary* action prescribed by that information-how leads to a failure to F . By contrast, an intelligent skilled capacity to F implies that the agent can improvise, will be disposed to supplant a prescribed step with a more efficient one and will not be disposed to repeat a mistaken step in doing F .

5. A Comparison with Previous Works

In this section, we discuss two similar accounts of information-how.

5.1. Control information and cybernetics

Peter Corning has argued that ‘control information’ is crucially missing in both Norbert Wiener’s cybernetic paradigm and the Mathematical Theory of Communication (Corning 2001). Control information is defined as the *capacity* to control the acquisition, disposition and utilisation of matter/energy in purposive processes. “If energy is [...] ‘the capacity to do work’, control information is ‘the capacity to control the capacity to do work’” (Corning 2001, p. 1279).

Control information, on Corning’s view, has three salient characteristics. First, it cannot be found in physical objects alone, but rather in a specific relationship between a given object and a user of that information. It is, thus, always relational and context-dependent. Second, until it is used, it only exists *in potentia*. An unread book of instructions or an undetected pheromone only represents potential control information (ibid). Third, although control information is a property of the relationships between “things”, rather than a physical entity, it *is* quantifiable. It can be measured in relation to what it does, that is, in relation to its “power” to control and utilise available energy and matter in or by a purposeful system (Corning 2001, p. 1280).

Corning’s characterisation of control information is certainly similar to that of information-how in some respects. On our view, too, information-how does not exist independently of a consumer. Rather, it exists as a dynamic relationship between an agent and the world (or another agent) or an agent and herself. There might be a seeming conflict between information-how depending on the *existence of an agent* and the *logical* relation information-how encompasses (as is discussed in Section 3.2). But this is no more problematic than information-how – in the logical realm – specifying *actions* that have *causal effects*. Information-how simply specifies actions that are to be performed by an agent

that can, in principle, perform them under the right conditions. (Consider information-how including the instruction ‘Jump’ for a person in a wheelchair.)¹⁵

What exists independently of the consumer is the physical *data*, which underlie and constrain the formation of possible information. The qualification of unused information-how as only being potential information needs unpacking. (Recall, according to PIH, it is the potentiality of semantic content to yield a particular result through action that makes it information-how.) It invites questions such as ‘Is information-how that the recipient *ignores* only potential information?’ and ‘Is information-how that is partially received by an agent leading to the particular result merely potential information?’.

Moreover, whether the best way to quantify information-how is by calculating its capacity to control and utilise available energy remains to be seen. Information-how should yield the same value for the utilisation of energy regardless of whether it is merely potential. At least *prima facie*, drawing on *data* as the vehicle of information opens up the possibility of quantifying both information-that and information-how using the theory of signalling games, which is widely used in biology.

5.2. *Instructional information*

On Floridi’s view, a piece of information is *instructional*, if it conveys the need for a particular action (2010, p. 34). Some instructional information is a subtype of semantic content, yet other is a subtype of what he calls ‘environmental information’ (Floridi 2011a, p. 92), “depending on whether *meaning* is a required feature” (Floridi 2010, p. 35, italics added). To see the distinction between the two types of “instructional information”, consider the difference between conventional Boolean gates and a dishwasher operation manual. According to Floridi, the former is an instance of *environmental* instructional information,

¹⁵ More fundamentally (as is mentioned in Section 3.1) this is related to the broader problem of meaningfulness of semantic content in the absence of an informee. Gestures or diagrams, for example, can also convey information-how but require *interpretation*.

since these gates “merely channel electric voltage[s] [...and] there is no semantics involved at the level of the gates” (2010, p. 35). The operations of these gates can, supposedly, be interpreted as IF/THEN logic instructions. On the other hand, the operation manual provides *semantic* instructional information on how to operate the dishwasher either imperatively (e.g., in the form of a recipe) or conditionally (i.e., in the form of an inferential procedure).

Floridi’s account of instructional information is not fully developed yet, standing in sharp contrast to his full-blown account of information-that (2011a, Chapters 4, 5). He does not specify any sufficient condition for either semantic content or environmental information to qualify as *instructional* information, rather than information-that, which is truth-evaluable. Some environmental information also qualifies as information-that, on his view (Floridi 2011a, p. 43).

Despite the similarity of our account to Floridi’s account of instructional information, they differ in many ways. On his view, some information, including instructional information as a subtype of *environmental* information, exists independently of any consumer. But, then, in what sense do conventional Boolean gates process instructional information? Floridi argues that they process instructional information as a subtype of environmental information. Yet, it is only when the input(s)/output(s) of these gates are interpreted by users that the channelling of electric voltages rises to the level of a *logic instruction*. Besides, it has been argued in Fresco and Wolf (2014) that most two-input, one-output Boolean gates do not even process “structured” data, let alone instructional information.

6. Conclusion

This article offers a new account of information-how that can, arguably, benefit the knowledge-how/knowledge-that distinction as well as have concrete applications in the

design of artificial agents. We have characterised information-how as a kind of semantic content that, unlike information-that, is not straightforwardly truth-evaluable.

When do data qualify as information-how? Data qualify as information-how when they rise to the level of semantic content whose fulfilment through action can systematically produce a particular outcome in a given context. Information-how has three notable features: (a) it is intimately action-directed relative to an informed recipient, (b) it can be effective or ineffective, and (c) it can be usable or not for the recipient. As such, for cognitive agents it expresses an expectation for some goal-directed action by the recipient in a given context.

Thus characterised, the more fundamental distinction between information-how and information-that can arguably shed light on the knowledge-how/knowledge-that distinction. Four important questions should be addressed to advance the knowledge-that/knowledge-how debate.

Q1. Does all information-how consist of information-that?

Q2. Is all information-how equivalent to information-that?

Q3. What is required for information-how to upgrade to knowledge-how?

Q4. What is required for information-that to upgrade to knowledge-that?

There are three main approaches to answering questions one (Q1) and two (Q2). On the strongest form of reductionism, *all* information-how consists of information-that. Thus, depending on the answers to questions three (Q3) and four (Q4), we may possibly learn that *all* knowledge-how, too, consists of knowledge-that. On the intermediate form of reductionism, whilst not all information-how consists of information-that, the former is *translatable* to the latter. They are *distinct* types of information, since the former consists of information-that plus some other thing(s). On the strongest form of antireductionism, the answers to Q1 and Q2 are negative. Any mere translation from information-that to information-how (and vice versa) loses (or adds new) information in the process. Thus, if

knowledge-that and knowledge-how are founded on information-that and information-how, respectively, they are distinct types of knowledge.

*Sidney M. Edelstein Centre,
The Hebrew University of Jerusalem*

References

- Billari, F. C., Fent, T., Prskawetz, A., & Scheffran, J. (2006). Agent-Based Computational Modelling: An Introduction. In F. C. Billari, T. Fent, A. Prskawetz, & J. Scheffran (Eds.), *Agent-Based Computational Modelling* (pp. 1–16). Heidelberg: Physica-Verlag.
- Cederman, L. E. (2001). Agent-based modeling in political science. *The Political Methodologist*, *10*(1), 16–22.
- Corning, P. A. (2001). “Control information”: The missing element in Norbert Wiener’s cybernetic paradigm? *Kybernetes*, *30*(9/10), 1272–1288.
doi:10.1108/EUM0000000006552
- Dretske, F. I. (1981). *Knowledge & the flow of information*. Cambridge, Mass.: MIT Press.
- Dreyfus, H. L. (2006). Overcoming the myth of the mental. *Topoi*, *25*(1-2), 43–49.
doi:10.1007/s11245-006-0006-1
- Fetzer, J. H. (2004). Information: Does it Have To Be True? *Minds and Machines*, *14*(2), 223–229. doi:10.1023/B:MIND.0000021682.61365.56
- Floridi, L. (2010). *Information: a very short introduction*. Oxford ; New York: Oxford University Press.
- Floridi, L. (2011a). *The philosophy of information*. Oxford: Oxford University Press.

- Floridi, L. (2011b). A Defence of Constructionism: Philosophy as Conceptual Engineering. *Metaphilosophy*, 42(3), 282–304. doi:10.1111/j.1467-9973.2011.01693.x
- Fresco, N., & Wolf, M. J. (2014). The instructional information processing account of digital computation. *Synthese*, 191(7), 1469–1492. doi:10.1007/s11229-013-0338-5
- Ghose, A. K., Morrison, E., & Gou, Y. (2013). A Novel Use of Big Data Analytics for Service Innovation Harvesting. In *Proceedings of the 2013 5th International Conference on Service Science and Innovation* (pp. 208–214). IEEE. doi:10.1109/ICSSI.2013.45
- Hamblin, C. L. (1987). *Imperatives*. New York, NY: Basil Blackwell.
- Hartland-Swann, J. (1956). The Logical Status of “Knowing That.” *Analysis*, 16(5), 111–115. doi:10.1093/analys/16.5.111
- Hartley, M., Fagard, J., Esseily, R., & Taylor, J. (2008). Observational Versus Trial and Error Effects in a Model of an Infant Learning Paradigm. In V. Kůrková, R. Neruda, & J. Koutník (Eds.), *Artificial Neural Networks - ICANN 2008* (Vol. 5164, pp. 277–289). Berlin, Heidelberg: Springer Berlin.
- Hetherington, S. (2011). *How to Know a Practicalist Conception of Knowledge*. Hoboken: John Wiley & Sons.
- Huttegger, S. M. (2007). Evolutionary Explanations of Indicatives and Imperatives. *Erkenntnis*, 66(3), 409–436. doi:10.1007/s10670-006-9022-1
- Jablonka, E. (2002). Information: Its Interpretation, Its Inheritance, and Its Sharing. *Philosophy of Science*, 69(4), 578–605. doi:10.1086/344621
- Kahneman, D. (2011). *Thinking, fast and slow*. New York: Farrar, Straus and Giroux.
- Lewis, D. K. (1969). *Convention: a philosophical study*. Harvard University Press.

- Meyer, J.-J. C. (2000). Dynamic Logic for Reasoning About Actions and Agents. In J. Minker (Ed.), *Logic-Based Artificial Intelligence* (pp. 281–311). Boston, MA: Springer US.
- Millikan, R. G. (2004). On Reading Signs: some differences between us and the others. In D. K. Oller & U. Griebel (Eds.), *Evolution of communication systems: a comparative approach* (pp. 15–30). Cambridge, Mass: MIT Press.
- Noe, A. (2005). Against intellectualism. *Analysis*, 65(4), 278–290.
doi:10.1093/analys/65.4.278
- Perry, J. (2003). *Knowledge, possibility and consciousness*. Cambridge, Mass.: MIT Press.
- Poston, T. (2009). Know how to be Gettiered? *Philosophy and Phenomenological Research*, 79(3), 743–747. doi:10.1111/j.1933-1592.2009.00301.x
- Reed, N., McLeod, P., & Dienes, Z. (2010). Implicit knowledge and motor skill: What people who know how to catch don't know. *Consciousness and Cognition*, 19(1), 63–76. doi:10.1016/j.concog.2009.07.006
- Ryle, G. (1949). *The concept of mind* (New University of Chicago Press ed.). Chicago: University of Chicago Press.
- Scarantino, A., & Piccinini, G. (2010). Information without truth. *Metaphilosophy*, 41(3), 313–330. doi:10.1111/j.1467-9973.2010.01632.x
- Silverman, E. (2007). *Methodological investigations in agent-based modelling*. University of Leeds, United Kingdom.
- Skyrms, B. (2010). *Signals: evolution, learning, & information*. Oxford: Oxford University Press.
- Stanley, J., & Williamson, T. (2001). Knowing How. *The Journal of Philosophy*, 98(8), 411–444. doi:10.2307/2678403
- Von Wright, G. H. (1951). Deontic Logic. *Mind*, 60(237), 1–15. doi:10.2307/2251395

Zollman, K. J. (2011). Separating Directives and Assertions Using Simple Signaling Games. *The Journal of philosophy*, 108(3), 158–169.