

Modeling Practical Thinking

Matthew Mosdell

1 Introduction

What is practical thought?¹ Traditionally, it has been distinguished from its theoretical counterpart by what it explains. Practical thinking explains the things we do; theoretical thinking explains the things we believe. Closely allied with this age-old distinction is a more recently emphasized one between knowledge-how and knowledge-that.² Again, following tradition, one knows *how* to ϕ when the object of knowledge is something the agent is able to do. I know how to ride a bike, bake a cake, or wiggle my ears, for example, if I'm able to do those things were I to try. When one knows *that* σ , on the other hand, the object of knowledge is propositional. I know that 'snow is white', 'diamonds scratch glass', or 'pens are mightier than swords' when those beliefs non-accidentally track the way things are. In either case, the objects of knowledge tend to be explained in terms of one or another of these distinctive ways of thinking.³ Theoretical thinking explains my knowing *that* 'the ground is white' when I see snow falling past my window; practical thinking explains my

¹Since 'thought' is ambiguous between act and object, and since it is the act I'm interested in, I'll use the slightly more cumbersome 'thinking' in this essay. What does the act of practical thinking, which is a movement from the objects of thought to intentional action, look like?

²The familiar distinction is made and defended by Gilbert Ryle in different places. See for example, Ryle (1946) and chapter 2 of Ryle (1949). The contemporary debates surrounding the distinction are nicely represented in Bengson and Moffet (2011).

³The phrase 'way of thinking' is something of a term of art. It means the manner in which information is presented to the mind. Terminological equivalents, which crop up in the literature about knowledge-how are 'mode of presentation', 'guise of thought', 'way of grasping', and 'practical sense'. Jason Stanley elaborates and defends a view of 'ways of thinking' that includes them as components of propositions in his (2011a) monograph. I'll have very little to say about Stanley's discussion, focusing instead on Carlotta Pavese's (2015b) notion of a 'practical sense'. See also footnote 13.

knowing *how* to move my legs in a circular motion when riding a bike. I'll assume, then, that knowledge-how stands to practical thinking as knowledge-that stands to theoretical thinking: the objects of knowledge are explained by their corresponding way of thinking. My interest will be exclusive to the practical side of things. Accordingly, when an individual knows how to ride a bike, bake a cake, or wiggle his ears, it's a distinctively practical way of thinking that explains what he knows when he knows how to do one of these things. But what makes a way of thinking distinctively practical?

There is an abundance of answers to that question,⁴ and I can't possibly hope to engage them all here. I do, however, want to focus on one fashionable contemporary answer. Intellectualists about knowledge-how have argued that an agent's ability to ϕ is explained by grasping a proposition (or set of propositions) that answers the question "How could one ϕ ?"⁵ But this view is inadequate unless the proposition that answers the question is grasped, entertained, or presented to the mind in a particular way. After all, one may grasp a true proposition that answers the question "How could one ride a bike" without thereby knowing how to ride a bike.⁶ What, then, is the way a proposition must be grasped in order for it to explain an agent's knowing how to ϕ ? The obvious (if trivial) answer is that one must grasp the proposition in a practical way. And this is the intellectualist response. An agent's knowing how to ϕ is explained by showing that a true description of how one could ϕ is

⁴Accounts of practical reasoning may seem tangential to the recent literature about knowledge-how, but I'm not sure that this perception is entirely justified. The puzzles, motivations, and interests share much in common. Discussions of rational agency, what it means for thought to give rise to action, and the role of practical knowledge in explaining the things we intentionally do seem to be a cluster of issues closely allied with those found in the recent know-how literature. See, for example, the work by Anscombe (2000), Bratman (2007), Davidson (1980), Frankfurt (1998), Korsgaard (2008), Rödl (2007), Setiya (2008), Thompson (2008b), Velleman (2000). It also seems to me that the literature discussing Elizabeth Anscombe's work is motivated by similar questions. See the essays in Ford et al. (2011) and O'Brien (2007), Rödl (2011), Schwenkler (2015), Small (2012), and Thompson (2008a).

⁵There are by now several versions of intellectualism. My characterization is in keeping with Stanley and Williamson (2001) and Stanley (2011a).

⁶See Schaffer (2007, 396) for the argument to support this intuition.

grasped, entertained, or presented to the mind in a practical way.⁷

Until recently, however, saying what it is to grasp a proposition in a peculiarly practical way has not been fleshed out. Prior to Carlotta Pavese’s “Practical Senses” (2015b), Jason Stanley had done the most to illuminate the idea, but his view came to little more than the thought that entertaining or grasping a proposition in a practical way is thinking of it practically.⁸ Pavese, in contrast, has attempted to illuminate the idea by giving an explication of practical modes of presentation in terms of Fregean senses. She has tried to say what it means to present the content of a proposition to the mind in a way that explains an agent’s ability to produce actions, and she has done so by interpreting Fregean senses as inferential rules that entail the ability to intentionally ϕ when facts about how to ϕ are fed through them as input. As a general explanation of what agents are able to do, however, Pavese’s strategy for modeling what it means to think practically falls short. My aim here is to show why.

In the next section, I frame the problem Pavese’s model of practical thinking was designed to solve and show how it works to do so. Section three articulates the presuppositions built into her account of practical thinking. Once I’ve ferreted them out, I’ll be in position to argue for the explanatory limits of Pavese’s model. This is what I do in section four by sketching three classes of action for which her model of practical thinking is a bad explanatory fit. If I’m right, the most promising intellectualist view of practical thinking to date provides at best a partial explanation of the things we know how to do. I conclude by considering some broader philosophical lessons of my arguments.

⁷The idea of a practical way of grasping a proposition or a practical mode of presentation was broached initially in Stanley and Williamson’s (2001) seminal paper. Critical responses followed quickly. See especially Koethe (2002) and Schiffer (2002).

⁸For Stanley’s discussion, see (Stanley, 2011a, 85-86). Pavese makes this point on page 2 of “Practical Senses” (2015b).

2 Practical Ways of Thinking

I will be using the notion of a model as it occurs in the philosophy of science to engage the intellectualist vision of the practical mind.⁹ Models are tools of mediation, bridging the gap between abstractions of the world and the world as it is. As such, they tend to introduce distortions that limit their explanatory scope. An early example of this is the mediating role played by Niels Bohr's hydrogen atom to explain the emission wavelength of hydrogen. The model helped to clarify puzzling features of hydrogen by showing how electrons moving between stable orbits could emit or absorb energy, but the model worked by distorting orbital structure and other spatial features of hydrogen. In the case at hand, Pavese has proposed a model that mediates between an abstract conception of practical thinking and practical thinking as it is in actual agents. Like Bohr's model, however, the one proposed by Pavese introduces a variety of distortions, which (I hope) will come into focus as we proceed. For that to happen, I need to bring the model itself into the open. Since Pavese's vision of practical thinking has grown out of intellectualist views of knowledge-how, I'll begin there.

2.1 Mind the Gap

According to intellectualists, knowledge-how is a species of knowledge-that.¹⁰ Or, to put the idea another way, knowing how to do something is knowing facts that truly describe how to do it. So, I know how to swim when I grasp a true description of how to swim in the appropriate way. Simple though the idea may seem, the devilish details are hidden in what it means to grasp facts about how to ϕ in *the appropriate way*. Making sense of

⁹See Downes (1992), Morrison (1999), Suppes (1977), and Weisberg (2013) for discussion of the role of models in scientific theorizing.

¹⁰Stanley and Williamson (2001), Stanley (2011a,b).

those details has been a long-standing challenge for intellectualists,¹¹ and it was initially met with a promise: the promise that facts, when possessed as knowledge-how, would not be behaviorally inert.¹² It was never exactly obvious, however, whether that promise could be fulfilled. Why would knowing a fact describing how one could ϕ entail the ability to ϕ ? In attempting to bridge the gap left by the not-so-obvious entailment, intellectualists have consistently appealed to practical modes of presentation, ways of thinking, guises of thought, and so on as the mechanism by which known truths avoid behavioral dormancy.¹³

The idea is simple: let the mechanism of a practical way of thinking (mode of presentation, practical sense, guise of thought) play a certain functional role in the mind of agents. When truths about how to ϕ are fed through the mechanism, actions are its output. In this way, information that truly describes how to swim has the effect of swimming when fed through a practical way of thinking. What is the nature of the mechanism? According to intellectualists, it is another sort of fact, which, along with descriptions of how to ϕ , is encoded in the content of propositions.

And just how is that supposed to work? In Stanley's (2011a) monograph *Knowing How*,

¹¹The intellectualist position was built upon, elaborated, and developed for approximately fifteen years before anyone took up the challenge, which was put forward clearly and powerfully very early on. See, for example, Koethe (2002) and Schiffer (2002).

¹²The promise is a gesture at an entailment relation: facts possessed in the right way entail an agent's ability to ϕ . This purported entailment is an important feature of the intellectualist view. Indeed, a variant of the commitment crops up rather frequently. Pavese, for example, is explicit that grasping a way of thinking (in her words, 'practical sense') endows an agent with certain set of abilities (Pavese, 2015b, 10). Bengson and Moffit (Bengson and Moffet, 2007) are also explicitly committed to ability-entailing concepts. Stanley is more guarded, though still committed. He writes, "the view that the contents of thought include ways of thinking ...entails that propositional knowledge is not behaviorally inert (Stanley, 2011a, 98).

¹³A note on terminology: as the main text indicates, the words used to discuss the topic of this essay vary. 'Practical mode of presentation' and 'practical way of thinking' are probably the most common, but Pavese trades those for the more Fregean flavored 'practical sense'. In each case, the terminology aims to capture that part of the mind that renders propositional content into an agent's ability to act. I drop the idiosyncratic language and instead speak merely of 'practical thinking'. Even so, the idea is always the same: practical senses, modes of presentation, ways of thinking, guises of thought, and so on all seem to refer to the same thing, which is practical thinking. See also footnote 3.

he argued that practical ways of thinking could be patterned after first-person ways of thinking. The idea was that since bringing *de se* facts before the mind requires a way of thinking those facts as facts about oneself, there is precedent for encoding ways of thinking into the content of thought. Here's the idea from Stanley:

The proposition that I am tired and the proposition that Jason Stanley is tired are distinct propositions, despite the fact that I am Jason Stanley. The former proposition contains a first-person way of thinking of Jason Stanley. My use of the first-person pronoun has a distinct propositional contribution from my use of the name "Jason Stanley"—the former expresses a first-person way of thinking, and the latter does not. (Stanley, 2011a, 104)

Stanley is noticing here that when thinking about oneself, the agent is required to adopt a way of thinking (or mode of presentation) that is itself part of that very thought. He uses this observation to suggest that practical ways of thinking might be similarly encoded as factual components of propositions, which means that ways of thinking are "constituents of propositions that we know" (Stanley, 2011a, 106). On such a view, grasping a proposition about how to do something—i.e., knowing how—entails grasping its practical mode of presentation, which, as with first person ways of thinking, is tied to an agent's power of action by the dispositions it presupposes. Stanley writes:

To think of an object in a first-person way is for that object to occupy a certain functional role—to be something towards which first-person dispositions are directed. Similarly, explaining what it is to think of a way of doing something in a practical way is . . . a matter of spelling out the distinctive practical functional role that way occupies in the mental life of the speaker. (Stanley, 2011a, 124)

But Stanley never moves past this gesture. He never spells out the "distinctive practical functional role" of practical ways of thinking. In the end, we're left wondering how the mechanism is supposed to work.

2.2 Modeling Practical Ways of Thinking

Pavese (2015b) remedies this problem by modeling practical ways of thinking (modes of presentation) as Fregean senses.¹⁴ And her view is thoroughly intellectualist. As she indicates early on, “practical modes of presentation are conceptual components of the propositional content that is putatively known when one knows how to do something” (Pavese, 2015b, 2). This means that when an agent knows how to ϕ , she knows a Fregean proposition composed of a description that truly answers the question ‘How could one ϕ ?’ and a conceptual rule for executing the descriptive input. An example or two will prove useful to show how propositions of this sort are supposed to explain knowledge-how.

Suppose I know how to add. On Pavese’s view, there are two components of the object of my knowledge. First, there’s the descriptive input that says, for example, “Take a number and *combine* it with another number to produce a third,” and second, there’s the inferential rule that specifies the sense of ‘combine’ given in the description. That is, the sense of the concept *combine* is specified by an inferential rule, and when information is presented under that concept, it generates a specific output. Another example: suppose I know how to make inferences in the manner of *modus ponens*. Again, the descriptive input might be something like, “Take an ‘If . . . , then . . . ’ statement and *combine* it with an assertion of the antecedent to produce the consequent as output.” The sense of that combinatorial procedure will be given by the inferential rule ‘modus ponens’. In both examples, the concept ‘combine’ serves the functional role of providing rules by which an executive module operates on input to produce output. Specifying the content of the ‘combine’ concept is tantamount to specifying a ‘way of thinking’ about the input.

¹⁴See Frege (1948) and Frege (1956) for the original characterization of a sense. Further discussion can be found in Burge (2005), Dummett (1993), and Evans (1982).

This manner of modeling ways of thinking should seem intuitively plausible. After all, many of us have thought about Turing machines and, even if we haven't, are familiar with the basic idea that computers perform operations on input to produce output. We might be fuzzy on the details of the operational procedures, but whatever they are, we're likely sympathetic to the idea that a few basic rules can be used to generate complex output. Similarly, even if we want to be agnostic about whether the operations performed by computers amount to thinking, we're probably sympathetic to the idea that such operations can *look* a lot like thinking. But never mind that. Whether computers think or not, the operational rules that govern the tasks they perform (the rules that determine what they know how to do) are used by Pavese to model the ways of thinking (modes of presentation, guises of thought) required for us to be capable of doing the things we do.

Let's suppose Pavese is right. When facts that answer the question "How could one ϕ ?" are grasped under a practical sense (an inferential rule that provides executive procedures), the result is the ability to execute an action up to a certain point (Pavese, 2015b, 13). Such abilities are entailed by an agent's grasp of a practical sense. Suppose I want to make cake and that I have the capacity to follow rules for doing so. In coming to grasp those rules, I thereby acquire the ability to make cake.¹⁵ The supposed entailment is even more obvious using a technological analogue. Suppose I want to run a Microsoft program on my Mac. To do so, I need to update my Mac's executive functions to accommodate the rules of the Microsoft program. Once the update has been made, once my Mac can execute the rules of the Microsoft program, it acquires a new ability. And this acquisition is a matter of entailment. Our minds must work in more or less the same way according to the model

¹⁵Pavese distinguishes between capacities and abilities. I'm following her in that distinction since it's a useful one. Accordingly, capacities are distinct from abilities in the sense that only the latter "requires that a subject possess a rule" (Pavese, 2015b, 10). I may have the capacity to bake cakes by following a rule, but only when I come to grasp the rule do I acquire the ability.

of practical thinking proposed by Pavese. Indeed, practical ways of thinking are modeled after the inferential rules of algorithms (which are instantiated on computers as programs). Consequently, coming to grasp a rule of execution—i.e., acquiring information encoded as a practical sense—is ability entailing.

This is an important development of the intellectualist view, since it's the clearest attempt to fill in the purported entailment between knowing propositional truths about ϕ -ing and knowing how to ϕ . Let me conclude this section, then, with one more example. Suppose I possess all the background capacities necessary to play the fiddle. I can hold the bow and fiddle in the right positions, saw the former across the latter in the relevant sort of way, tap my feet to the strum of the banjo, whatever. “Were [I] instructed to ϕ in a certain way, [I] would be able to ϕ *by following that instruction*” (Pavese, 2015b, 9-10, emphasis in original). Capable though I may be, I don't know how to fiddle. I don't have the ability. So, on Pavese's account, what do I need to learn in order to acquire the ability? First, I must come to understand information describing a way I could play the fiddle, and, second, I must grasp an inferential rule that when executed would produce the act of fiddle-playing.¹⁶ It is the latter that renders information describing how to play the fiddle into my playing the fiddle. Once I've got the information and the rule, I'm done. The gap between knowing a true description of how to ϕ and the ability to ϕ is closed.¹⁷ As Pavese writes, “by grasping a practical sense [practical mode of presentation or practical way of thinking], one is endowed

¹⁶And let me again emphasize that these rules are encoded as practical senses, which are conceptual components of Fregean propositions known by an agent when she knows how to ϕ . It's in this sense that “knowing a rule is a matter of possessing a practical concept (Pavese, 2015a, 22).

¹⁷It's worth emphasizing that the ability to ϕ and being able to ϕ may still come apart on Pavese's model (Pavese, 2015b, 11). I may grasp a rule for playing the fiddle and know that *this* is a way to play the fiddle, which means I have the ability, and still be unable to do so, since I recently lost both of my arms to a crazed lion.

with a certain set of *distinctive abilities*” (Pavese, 2015b, 10, emphasis in original).¹⁸ Keeping with the example: by grasping a rule for fiddling—that is, by bringing an algorithm for fiddling before my mind—I’m endowed with the ability.

Let’s recap: Intellectualist accounts of knowledge-how have attempted to bridge the gap between what we know and what we are able to do by appealing to a peculiarly practical mode of presentation or way of thinking. The bridge is necessary in order to explain how we get from knowing that some description of ϕ -ing is a way to ϕ to the ability to execute that way of ϕ -ing in action. Characteristic of such accounts is the idea that a component of the propositional content encoded in states of knowledge-how entails the ability to do this or that. On the most substantive version of this idea, the one we’ve just finished sketching, practical ways of thinking are modeled as inferential rules for executing commands describing a way of doing something. Consequently, an agent’s knowing that this is a way to ϕ endows her with the ability to ϕ when she grasps both that *this is a way to ϕ* and *the rules necessary for executing that way of ϕ -ing*. Both components are encoded in the propositional content of states of knowledge-how and *entail* the ability to ϕ when grasped.

3 Explanatory Expectations

If inferential rules are supposed to explain how we get from static propositional content to an ability to execute it in action, we should wonder how, exactly, that explanation works. On Pavese’s model, the explanation is one of entailment: an agent’s grasp of the inferential rules characteristic of a practical sense entails the ability of that agent to ϕ . But does it? How? What’s the nature of that entailment? The weight of the answers to these questions will have

¹⁸As Pavese notes, she’s following Bengson and Moffet (2007) here. In contrast, however, her view is designed to apply to all instances of know how and is explicitly Fregean.

to be born by the similarity between the rules by which computers execute algorithms and our minds, for it is upon that similarity that the plausibility of Pavese's model of practical thinking turns. If the similarity is as robust as Pavese believes, we'll learn a lot about our ability to produce actions by reflecting on computer algorithms. Dissimilarities, however, will set explanatory limits.

How can we determine how robust the similarity is? Since we don't actually know what practical thinking looks like, we can't determine how similar it is to the algorithms used by computers without begging questions. Even so, the idea that an agent's ability to ϕ is entailed by her grasp of inferential rules for ϕ -ing provides an avenue of comparison. Let me explain. Since certain features of algorithms are required in order to be implemented by computers, and since Pavese model's practical thinking after computer algorithms, we can get a clearer picture of practical thinking by making the features necessary for an algorithm's execution by a computer explicit. By understanding the entailment relation between a computer's acquisition of an algorithm for ψ -ing and its ability to ψ , we'll be in a position to understand the purported entailment between an agent's grasp of the inferential rules for ϕ -ing and his or her ability to ϕ . Why? Because the weight of that purported entailment is born by the similarity between the workings of a computer and the workings of an agent's mind.

What must an algorithm look like if its possession is going to entail a computer's ability to execute it? I want to consider three features. First, algorithms must have parts: they must be composed of well-defined operations that terminate in well-defined solutions. Second, algorithms must be specified in a formal language that is unambiguous and precise. And third, algorithms are wholes that are designed to accomplish a specifiable solution.¹⁹ Given

¹⁹The features of algorithms I'm emphasizing are culled from standard definitions found in computer science textbooks. For example, J. Glenn Brookshear's widely used *Computer Science: an overview* defines

these features of algorithms, we can articulate a trio of features presupposed by Pavese’s model of practical thinking. Practical ways of thinking should be 1) composed of well-defined operations, 2) formally specified in a way that is unambiguous and precise, and 3) designed to achieve a specifiable solution.²⁰ Again, a reminder is in order. Practical ways of thinking (modes of presentation, practical senses) are “conceptual components of the propositional content that is putatively known when one knows how to do something” (Pavese, 2015b, 2). And when one knows how to ϕ , one knows “a practical proposition, a proposition that has as a component a practical sense” (Pavese, 2015b, 2). This means that knowing how to ϕ involves grasping a concept with the features I’ve just specified. Only when such a concept is grasped will the gap between what an agent knows and what she is able to do be bridged in a satisfactory way. But how far can such a model actually take us in filling the explanatory gap between knowing a true description of how to do something and being able to do it?

an algorithm as “an ordered set of unambiguous, executable steps that defines a terminating process” (2014, 213). Similarly, Allen Tucker’s *Computer Science Handbook* defines an algorithm as “a finite sequence of instructions that is supposed to solve a particular problem” (2004, 161).

There are lots of things about algorithms that I’m happy to ignore, since those features won’t help me to make my argument. For example, algorithms are typically assumed to be general, finite, unique, abstract, coded using some specific language, mind-independent, and so on. These assumptions may or may not present potential problems for Pavese’s model, but whether they do or not, I’m inclined to ignore them. Similarly, there are a host of presuppositions involved with the execution of algorithms on specific computers. Is the computer on? Does it have a virus? Are all its parts functioning properly? Again, these questions are irrelevant to my aims.

²⁰A trio of comments corresponding to each of the above is in order. First, Pavese is explicit that a practical sense is an inferential rule with “part-whole structure such that every of its parts is identical to or is reducible to primitive inferential rules” (Pavese, 2015b, 13). Second, there is an attempt by Pavese to move away from the syntactical rigors of computer programs in order to allow practical senses to be something other than linguistic representations. She writes, “the inputs and outputs of practical senses are *ways of representing the commands to be executed and the result of the execution up to a certain point*. Such representations do not need to be sentences, and may be more picture-like or map-like” (Pavese, 2015b, 13). Since the strength of her view depends on the analogy between the functioning of computers and our minds, I believe this move is disingenuous. Even so, I’m happy to allow that practical senses may be map-like or picture-like, so long as those maps or pictures are formally precise, unambiguous, and discriminable. Finally, the third feature I mention—namely, that practical ways are designed to achieve a specifiable solution—is not meant to explain the intentionality of know how. For Pavese, that explanation is had via propositional knowledge, a component of which is one or another practical way (Pavese, 2015b, 17). My point is rather that as a concept representing a way to ϕ and the task of ϕ -ing up to a certain point, practical ways of thinking depend on ϕ being well-specified—i.e., what falls under the scope of ϕ cannot be ambiguous, vague, indeterminate, or indiscriminable.

4 Expectations Unmet

There are limits to what Pavese's model can represent and what inferences it is reasonable to draw from it. This is true of any model. Plato used the structure of a city to model a just soul, but to expect the soul to *actually* be structured in the manner of Plato's imagined city is absurd. If we're not mindful of that fact, we're bound to draw absurd conclusions. So, too, with Bohr's model of the atom: since it misrepresents its target in a variety of ways, drawing inferences from it must be done cautiously. The point applies to Pavese's account of practical thinking. Even if her model accurately represents the relation between knowing how to ϕ and being able to ϕ in some cases, its explanatory scope will be bound by the limitations of the model. In this section, I'm going to pursue those limits, but before I do, I want to get a sense of the area of explanation for which her model is a good fit.

Since Pavese models practical thinking after computer algorithms, her account should nicely represent actions performed by executing a sequence of instructions. Tasks like performing addition, modus ponens, multiplication, or Fibonacci number generation, which are all Turing computable tasks, find a nice fit in Pavese's model. Similarly, her model nicely extends to other rule-governed abilities. You want an explanation of your ability to make Indian curry, to do the hokey-pokey, or to assemble IKEA furniture? Look to a model designed after the sequenced commands found in computer algorithms. And there's little reason to think we couldn't construct formal algorithms with well-defined steps and clear termination points to represent a whole variety of abilities. For all of them, Pavese's model will be a nice fit.

There are, however, a host of abilities that will not sit well with the model of practical thinking put forward by Pavese. If the thinking that goes into executing an action lacks part-whole structure, is imprecise, or without a specifiable solution, Pavese's model will not

be able to represent it without introducing costly distortions.

4.1 Basic Actions

Following Arthur Danto (1965), it is standard to distinguish basic from complex actions.²¹

As Danto explains, “when [an agent] performs a basic action, he does nothing first that causes it to happen.” Basic actions are “perfectly simple in the same sense in which the old ‘simple ideas’ were said to be: they were not compounded out of anything more elementary than themselves, but were instead the ultimately simple elements out of which other ideas were compounded” (Danto, 1965, 147). As with ‘simple ideas’, basic actions are explanatorily primary.²² Nothing more fundamental explains our ability to perform them. And this is so by design: basic actions are required by theories of action to stop an impending regress. There have to be things agents do that are not done by doing anything more basic. Such doings, such abilities, are where explanation stops. We can identify them in a theory by reaching for an explanation—that is, by asking “How does one ϕ ? So, for example, “How do you do the hokie pokie?” reveals that the hokie pokie is not basic or elementary. In contrast, wiggling your ear might be. How do you wiggle your ear? How do you extend your index

²¹Doug Lavin Lavin (2012) describes the pervasiveness of basic action when he writes, “The classification [of actions into the basic and the non-basic] is meant to be one we must recognize if we are to understand the very structure of intentionally doing something: whatever large-scale projects one has realized through the ordering of means to ends, one must eventually reach a fine enough resolution and come upon things that have been done without any thought about how to get them done. This bare-bones depiction of the structure of practical or productive consciousness is not meant to be, in the first place, the upshot of theoretical investigation, but instead part of the pre-theoretical scaffolding on which our researches into the nature of action take shape: its apparently unproblematic inevitability rests on its being placed among the innocent preliminaries. The model is pervasive.” In arguing against the apparently unproblematic inevitability of basic action, Lavin is taking on a tradition whose influence is hard to overstate. Ever since Donald Davidson’s (Davidson, 1980) “Actions, Reasons, and Causes,” what has come to be called the standard story (Velleman (2000); Smith (2004); Hornsby (2004)) of action and agency has been taken to be a starting point for much theorizing. Everyone who accepts that starting point is committed to basic actions in the sense embraced by Danto. For influential examples, see Bratman (2007); Frankfurt (1998); Searle (1983); Setiya (2007); Velleman (2000), and others.

²²Arguments for moving away from the primacy of basic action can be found in Lavin (2012), Small (2012), and Thompson (2008b).

finger, raise your arm, wink, or wiggle your toes? If the performance of these things is basic, they will be performed by the agent directly and not by the agent's performance of anything more simple. We can describe this feature of such abilities (whatever they turn out to be) by saying they are structureless: there are no parts to their performance that *the agent* can perform.

In Pavese (2015b), she offers a conception of basic actions at odds with the one just sketched. According to her, “what makes an action basic (for a person) is that it is executable by that person by following a primitive and basic rule”²³ (Pavese, 2015b, 14). This is an unfortunate appropriation of terminology, but it needn't get in the way of understanding Pavese's view. Consider the act of addition: according to Pavese, addition is a basic action because it can be executed by following a primitive and basic rule. The **R-ADD** command is found in most programming languages for moving between and updating registers according to the rules of addition. Following it, a computer is able to perform one of Pavese's basic actions. Notice, however, that **R-ADD** works by a computer's performance of more basic operations: first, the computer reads the program text; second, it interprets variables; next, the computer updates its store; and so on. The tension with Danto's notion of basic actions is obvious. Since there are things that must be done by the computer in order to execute one of Pavese's basic actions, those actions are not basic in Danto's sense. Instead, the more basic operations of reading the program text, interpreting variables, etc., which are executed in the service of **R-ADD**, are basic in the relevant sense.

A role similar to Danto's basic actions is played in Fodor's work by ‘elementary operations’. According to Fodor, “an elementary operation is one which the normal nervous

²³Pavese tells us that primitive rules are not equivalent to any sequence of other rules, and basic rules do not contain any other rules as parts. See (Pavese, 2015b, 5-6) for several examples of rules that are primitive, basic, or both.

system can perform but of which it cannot perform a proper part” they “. . . are those which have no theoretically relevant internal structure” (Fodor, 1968, 629). As with Danto’s basic actions, Fodor’s elementary operations are designed to stop an explanatory regress. They are operations executed without doing anything more basic. In contrast to Danto’s view, however, elementary operations are sub-personal. They are not performed *by the agent*, but by one of its sub-systems. Pavese embraces the notion of elementary operations as “theoretical posits that are needed if we are to make sense of the possibility of computational explanations of behavior” (Pavese, 2017, forthcoming), but seems to distance herself from Fodor’s commitment to their role as *sub-personal* systems. Indeed, in her discussion of practical senses (Pavese, 2015b), she says she is “appealing to practical senses to provide a personal explanation of a certain subset of a subject’s abilities” and “the level of explanation at which [her] arguments work is personal: . . . [her] abilities only encompass competences that a *subject* can be properly said to possess” (Pavese, 2015b, 18, emphasis in original). I suspect this focus is why she says “an elementary operation for a system *s* at a time *t* is one that *s* can perform at *t* but of which *s* cannot perform at *t* a proper part” (Pavese, 2017).²⁴ If that’s right, then Pavese’s account of practical thinking explains an agent’s ability to ϕ by primitive rules of inference, which are themselves composed of more simple, elementary operations *the agent* is able to perform (Pavese, 2015b, 2017, forthcoming). Alternatively, we might want to interpret elementary operations as Fodor does, as sub-personal performances. On this interpretation, the primitive rules of inference used by Pavese to explain an agent’s ability to ϕ bottom out in operations the agent *does not* execute or control.

²⁴I read this description of elementary operations as saying that the system/agent/person is in control of them. If that’s the right reading, Pavese is here conflating her ‘elementary operations’ with Danto’s basic actions. There are other ways of interpreting her, however. Following Fodor, she may mean that elementary operations are sub-personal and not controlled by the system/agent/person. I take up both alternatives in the main text.

On either option, Pavese's model faces difficulties when explaining basic actions (in Danto's sense). Suppose first that an agent's ability to perform basic actions is explained by rules of inference, which are themselves composed of elementary operations the agent is able to perform. Is the performance of such operations explained by a rule of inference? If so, we've embarked on a regress. If not, Pavese's model of explanation doesn't account for the agent's ability to perform them. Alternatively, suppose an agent's ability to perform basic actions is explained by rules of inference, which are themselves composed of sub-personal elementary operations the agent doesn't perform. This option has the virtue of avoiding the potential regress, but it still doesn't help Pavese. Positing sub-personal performances in an effort to explain an agent's ability to ϕ isn't to give an explanation of those posits. It is, rather, to relocate the explanatory problem. But even if we were to allow this move, there is the further worry that it leaves Pavese face-to-face with the unenviable task of trying to say how actions executed and controlled by the agent emerge from sub-personal operations that are not executed or controlled by the agent.

All this means that at some most basic level, Pavese's computational model cannot explain an agent's ability to ϕ . This problem arises from the very nature of her model: since the model's explanatory power relies on mapping practical ways of thinking to algorithms, abilities that are not the product of algorithmic functioning go unexplained. How pervasive are such basic/elementary abilities? Who knows. But there seems to be a range of things agents are able to do that are genuinely basic. For example, the ability to wiggle an ear, open your mouth, lift an arm, look over there, bend your toes, turn your head, entertain a proposition, and so on seem to be basic in the relevant sense. But I'm not committed to this list and I'll not argue over the particulars. Whatever abilities Pavese wants to accept as basic, they'll go unexplained by a model that relies on practical modes of presentation,

practical senses, or practical ways of thinking, which require more basic abilities as parts.

Perhaps this limitation of the model can be ignored. Even if practical thinking can't explain basic abilities, there will be a range of unsophisticated abilities—e.g., the ability to add, make inferences using modus ponens, or sort objects into groups—that do get explained by the type of rules characteristic of Pavese's model. Whatever that range of abilities turns out to be, it will be a useful starting point for a project that promises substantial theoretical advantage. The model may fail to explain the ability to perform basic actions/elementary operations, but it certainly explains *some* class of abilities. Whatever that class is, it will serve as the basis for a project that shows significant promise. Perhaps, then, we should just ignore the model's limits, bite the bullet, and move on.

While I'm sympathetic to moving forward in this way, the promise of an account such as Pavese's isn't a compelling reason to buy into it as a general view of practical thinking. This observation won't come into focus until we've traversed sections 4.2 and 4.3, but in the meantime we should ask why Pavese might be willing to abandon explaining our ability to perform basic actions/elementary operations? Ignoring the model's explanatory limits might seem appealing. What advantage of Pavese's model is significant enough to ignore its explanatory limits? According to Pavese, the "most important advantage of [her] proposal . . . is that it affords a desirable account of how it is that the ability to perform a complex task *arises* from the ability to perform more basic tasks that are the parts of that complex task" (Pavese, 2015b, 15, emphasis in original). In other words, the chief advantage of her model is that it composes. It should be clear how this compositionality is supposed to work. Since practical thinking, or grasping a practical sense, is modeled after computational algorithms, the rules that constitute practical thinking should conjoin to produce more complex instances. Just as different inference patterns—e.g., modus ponens, modus tollens,

disjunctive dilemma—combine to produce more complex inferential chains, and just as simple concepts combine to produce complex sentences, the rules of practical thinking combine to produce more complex thoughts. If my ability to break eggs is explained by appealing to a simple rule of inference, and if my ability to whisk eggs is similarly explained, then these rules will be part of the combined explanation of my ability to make an omelette. The whole is explained by its parts.

As I've mentioned, I'm happy to concede this advantage of her model: our ability to perform algorithmic-like actions—i.e., actions with clear, precise steps terminating in a well-defined solution—might be best explained using an account of practical thinking like Pavese's. It seems perfectly plausible that we *could* construct algorithms to represent the practical thinking needed to execute a variety of abilities. And in cases of this ilk, the compositionality of Pavese's model does seem an asset. But even so, her model relies on a set of presuppositions that will not do well in accounting for other of our abilities. Of particular concern in this section has been the model's incapacity to explain our ability to perform basic actions/elementary operations. That's a significant cost to her model. The most basic/elementary things we are able to do go unexplained.

4.2 Abilities Relying on Analogue Magnitude States

Recently, Jacob Beck (2012) has argued for a class of mental states with nonconceptual content. He labels this class of mental states 'analogue magnitude states' and shows that they are characterized by "a systematic limitation, *Weber's Law*, which holds that the ability to discriminate two magnitudes is a function of their ratio" (Beck, 2012, 569). For example, suppose I'm standing in a courtyard with two buildings on either side of me and I'm asked to determine how far the entrance to each building is given my current location. As the ratio

of those comparative distances approaches one, my ability to determine which entrance is farther away diminishes. Similar examples apply to judgements of duration, rate, force, and so on. A thought experiment will help make the notion of an analogue mental state clear.

Imagine that you keep track of the number of people in a room by filling a bucket with a hose. Every time someone enters the room, you turn the hose on for about a second; and every time someone leaves, you pour a little bit out of the bucket. The height of water in the bucket will then be a direct analogue of the number of people in the room, and so you can use it as a decent approximation of that value. Of course, the representations provided by this analogue bucket system will not be perfectly precise. Given the imprecision in your method, your bucket representations will be intrinsically noisy. Thus, if you have two buckets representing the number of people in each of two separate rooms, your ability to reliably discern which room has more people will be a function of the ratio of the number of people in each room. As the ratio approaches one, the relative heights of the two buckets will become decreasingly reliable indicators of which room has more people, and below a certain threshold they will not be reliable at all. (Beck, 2012, 590)

It doesn't matter for our purposes whether the content of analogue mental states qualifies as conceptual. It does matter, however, that their content is inherently noisy, resulting in judgements of magnitude (size, duration, rate, force, etc.) that are comparatively imprecise—that is, judgements that are comparatively indiscriminate, vague, or ambiguous.

Before turning to the argument, some preliminary observations are in order. First, I'm taking the success of Beck's arguments for granted and assuming that some mental states are genuinely analogue. More particularly, I'm assuming that judgements of magnitude become indiscriminate as their comparative ratio approaches one. Given this assumption, we should expect that representing such states using models that rely on precision will inherently distort their analogue features. Second, I'm assuming that the target of Pavese's model is human thinking. The model seeks to explain how human agents move from propositional knowledge about ϕ -ing to ϕ 's execution in action. Key to her explanation of that movement are rules of inference modeled after the formally precise commands of computer algorithms.

If the analogy between the workings of a computer and the workings of the human mind breaks down, the model loses its explanatory power, since the supposed entailment relation that explains how we get from knowing that ϕ to an ability to ϕ is undermined. Of course, on the assumption that Beck's arguments are successful, the analogy is broken. Consequently, Pavese's model isn't positioned to explain the relation between an agent's use of analogue states in thought and the abilities tied to those thoughts. And I should note, whether computer software can be engineered to mimic abilities arising from analogue states is irrelevant. The mimicry, even if effective at producing actions that resemble human actions, depends on precisely specified rules of execution, which are incongruous with thoughts that rely on analogue states. With these points in mind, the success of this section depends on showing that analogue judgements are used in the practical thinking of human agents. Are they?

Consider an example or two. Suppose I'm a weekend mechanic who enjoys spending time rebuilding automotive engines. An important part of that activity is the ability to tighten bolts, which we might try to explain using a model of practical thinking that relies on precisely specified rules of execution. According to such a model, the mechanic's thinking would be best described by a rule like, "Tighten the exhaust header bolts to a torque of 15 newton meters" or "Tighten the bolts while torque is less than 15 newton meters." But such rules are implausibly precise. A better, more accurate characterization of the weekend mechanic's thinking would be something like, "Tighten the bolts 'til they're snug." Why is this a better characterization of the rule? Because human agents can't distinguish bolts tightened to 14 rather than 15 newton meters. The agent's judgment of force is represented with an analogue state. Another example: suppose my weekend hobby is target shooting. An important part of this hobby is to make adjustments to one's aim in order to hit targets at different distances and under different wind conditions. We could try to model that ability by

appealing to a rule whose content is precise, “For every 10 meters of distance, adjust the aim upwards by .25 centimeters; and for every 1 mile per hour increase in windspeed, adjust the aim in the opposite direction by .5 centimeters.” But again, the rule is an implausibly precise description of the target shooter’s practical thinking. A more accurate representation would appeal to a rule whose content relies on analogue magnitude states. Why? Because as targets move (or as the shooter moves to different targets), comparative judgements of magnitude (i.e., windspeed and distance) will be imprecise—that is, they will be indiscriminate, vague, or ambiguous. It wouldn’t be difficult to extend the list of examples indefinitely. Many of our abilities require making judgements that involve magnitudes of distance, rate, duration, force, and so on. (A few examples: driving a car, adjusting the volume of one’s voice in a loud bar, playing golf, figuring out when to leave the house in order to make the appointment on time, monitoring one’s jogging speed, and so on.) Many of those judgements will become indiscriminate as their comparative ratio approaches one—that is, those judgements will best be represented using imprecise analogue magnitude states. That shouldn’t be surprising given 1) that we have evolved from creatures that lack robustly discriminate representational capacities and 2) that much of what we do requires quick judgements of one or another type of magnitude. Consequently, a model of practical thinking that assumes inferential rules built from precise representational components will inherently distort the analogue features of human thought.

My opponent might concede the general point but worry that my examples are no good, since autonomous robots can perform all the tasks I’ve just described: they can be engineered to tighten bolts, measure windspeed and distance, play golf, drive a car, and so on. Consequently, the computational rules of execution designed to perform these tasks are sufficient to account for the corresponding abilities. That objection, however, misses its mark.

My point isn't that practical thinking in cases like those I've just mentioned couldn't be modeled using an account like Pavese's. Of course we can represent thoughts that make use of analogue states with models that are formally precise. But doing so assumes a view of the mind that is at odds with research like Beck's—research that resists the urge to see the mind as a kind of digital computer—and produces a distorted representation of the mind's functioning. In assuming the veracity of Beck's view, I'm assuming that analogue judgements—those that are comparatively indiscriminate, vague, or ambiguous—are part of the human mind, a part needed to explain abilities like those described above. But surely, my opponent might continue, the fact that robots can be engineered to perform tasks like those described shows the error of adopting views like Beck's. Hardly. Nothing about our own cognition follows from the fact that a robot or computer can be engineered to execute a task using precise rules.²⁵

Indeed, the fact that computers are the product of engineering weakens the plausibility of arguments that rest on the analogy between the workings of a computer and the workings of a human mind. That the various input modules of computers are engineered to provide information in a form that meets the precise specifications of the operational rules used to execute specific tasks contrasts with our own input and processing media. In our case, there is no fortuitous correlation between the things we know how to do and the information acquired through our senses. Our sensory apparatuses aren't *designed* for specific tasks; rather, they are *evolved* general-use mechanisms that help us to manage our way through an unpredictable world. Unlike computers (or the robots that operate on computer software), we have evolved sensory instruments that measure magnitudes in a rather imprecise way, so we should expect thinking that makes use of those instruments to be similarly imprecise. In

²⁵That is, nothing follows except the fact that human cognition is capable of engineering robots to operate on precise rules.

fact, getting human thinking to align with the precision typical of formal models like Pavese's requires inventing tools or instruments that can provide it. That is, precision thinking is an engineering accomplishment, not a naturally evolved feature of human cognition,²⁶ so we shouldn't expect our stand-alone thoughts to be precise in the way demanded of computer algorithms.

Let me flesh this out a bit more using examples from above. I've just pointed out that the design specifications of robots include measuring instruments that provide appropriately precise information to the software upon which the robot operates. So, for example, a robot designed to tighten header bolts comes equipped with an instrument for registering torque. Similarly, robots designed to autonomously measure magnitudes of distance and windspeed are equipped with the right instruments. More generally, equipment providing input will be designed to register measurements in increments corresponding to the precision of the rules governing the robot's execution of various tasks. There are no analogous task-specific instruments for human beings. The weekend mechanic doesn't possess a built-in instrument for measuring torque to meet the specifications of the various tasks he hopes to perform, so it would be odd if the thoughts required for performing those tasks relied on algorithms that demanded one. Similarly, the target shooter doesn't come equipped with instruments for discriminating closely associated magnitudes of distance and windspeed, so we shouldn't expect the practical thinking that goes into performing his hobby to operate on rules that require them. Of course, if the auto mechanic wants bolts tightened to a specific force, she could purchase a torque wrench. And if the target shooter wants a more precise accounting of changes in distance and windspeed, he could purchase a range finder, an anemometer, and a scope. The very fact that these tools must be engineered and purchased in order to eliminate

²⁶For a compelling argument to the effect that sharpening up our thinking and eliminating approximation and vagueness requires engineering, see Elijah Millgram's *Hard Truths* (2009).

imprecision in our sensory input, however, is evidence that our practical thinking in such cases is not naturally precise in the way demanded of rules governing a robot’s performance of these tasks. More generally, whenever comparative judgements of magnitude are noisy in a way that makes them imprecise, we should expect the practical thinking that relies on them to be poorly modeled with formally precise algorithms. In other words, a model like Pavese’s, which aims to represent practical thinking using rules modeled after formally precise algorithms, can only represent thinking that relies on analogue states by distorting it.

Such distortions come with costs. In cases like those I’m focused on here—i.e. cases where practical thinking makes use of indiscriminate, vague, or ambiguous mental states—the distorting effect undermines what Pavese considers the principal advantage of her model. Recall that according to Pavese, the promise of compositionality²⁷ is her model’s chief advantage. But if that model represents practical thinking that relies on analogue magnitude states by distorting it, her model’s chief advantage is actually a liability. If modeling a target shooter’s practical thinking with precise rules misrepresents what’s going on in his mind, this misrepresentation is compounded when his thinking involves serialized steps that make use of analogue magnitude judgements. Suppose, for example, that our weekend marksman is pretty good at discriminating targets that are approximately 100 yards away. In such cases, the distorting effects of a model like Pavese’s are inconspicuous and innocuous. Suppose further, however, that the marksman’s judgements of distance are increasingly noisy as his attention moves to targets farther downrange. That is, suppose the the noise of his analogue judgments aggregates as he moves to discriminate the distance of downrange targets. Without specialized equipment to manage his judgements, this should be expected.

²⁷Compositionality is the idea that “the ability to perform a complex task *arises* from the ability to perform more basic tasks that are the parts of that complex task” (Pavese, 2015b, 15, emphasis in original).

The growing imprecision of his judgements of distance will inevitably infect his practical thinking, which means that the marksman's ability will be increasingly misrepresented by a model like Pavese's. What were minor, innocuous distortions in the model's representation of the target shooter's ability to hit targets at 100 yards accumulate and serve to produce a thoroughgoing misrepresentation of the marksman's practical thinking when trying to hit targets farther downrange.²⁸ The promise of compositionality, which Pavese sees as the principle advantage of her view, becomes a liability. The composition of formally precise rules increasingly misrepresents the practical thinking of agents using analogue magnitude states to perform a task.

To recap: if the gap between knowing how to ϕ and being able to ϕ is explained by a way of thinking that is governed by rules of execution, we should wonder what those rules look like. On Pavese's model, they must be formally precise with unambiguous initiation and termination points. Only if such conditions are presupposed will there be anything like an entailment relation between grasping an inferential rule for ϕ -ing and an agent's ability to ϕ . But as we've just seen, many of our abilities aren't best explained by rules of this sort. Instead, rules that rely on analogue magnitude states better explain a range of abilities. In such cases, I may know that "Adjusting my aim for windspeed and distance" is how to hit the target, but since the rule involves analogue magnitude states, my knowing that ϕ doesn't entail my ability to ϕ and a model that relies on that entailment gets things wrong.

²⁸There is a problem here in the neighborhood of the sorites paradox. Recall that by chaining a bunch of seemingly benign instances of modus ponens together, one can be led from the true belief that 'Yul Brynner is bald' to the false belief that 'Jerry Garcia is bald' (Shapiro, 2008). After all, one hair never makes a difference. I don't want to take up issues of vagueness here, but I do want to point out that Pavese's commitment to formally precise rules shouldn't be surprising given her intellectual heritage. Timothy Williamson isn't merely the father of contemporary intellectualism, he's also one of a handful of epistemicists. See Williamson (1994, 2003, 2002) for examples of his epistemicist view.

4.3 Defining a Solution

An agent's grasp of rules for executing an action, ϕ , is designed to fill the space between knowing that ϕ and the ability to ϕ . Grasping such rules is grasping a practical sense, way of thinking, or mode of presentation. It is practical thinking. But must practical thinking be rule-governed if it is to fill that gap? I don't believe so, but I haven't tried to argue for that view here. I have, however, argued that some practical thinking is poorly modeled using rules structured after the manner of computer algorithms. Some of our abilities—namely, those to perform basic actions/elementary operations or actions produced from analogue magnitude states—are poorly modeled with structured, formally precise rules. I now want to turn to a different problem for such models by showing that our ability to discover, innovate, and invent solutions is also not well modeled with Pavese's account.

Algorithms are well-defined wholes only when they are composed of clear steps that terminate in a well-defined solution. It's that last feature that interests me in this final stretch. That algorithms terminate in well-defined solutions means that a model of practical thinking designed after them will face difficulties representing abilities that lack such termination points. I'm not by this suggesting that the variable that constitutes an algorithm's solution-space needs to be known prior to executing the algorithm; rather, I'm suggesting that the *type* of solution must be known prior to the algorithm's execution. There must be clear, unambiguous conditions delimiting the scope of the solution space. For example, you can't build a computer to solve addition problems if you don't know that the solutions must be numerical. Similarly, building a babysitting algorithm demands a clear conception of the range of solutions to such a problem.²⁹ What steps when followed as rules would effect successfully babysitting a three year old? If you can't envision and specify a range of success

²⁹These two examples (and their textual proximity) are taken from Pavese (Pavese, 2015b, 2).

conditions, the task of specifying the rules to achieve it is impossible.³⁰

These are points about a structural feature of algorithms. Their design is top-down.³¹ As a result, you have to know the kind of solution you're looking for prior to building an algorithm that eventuates in that kind of answer. But this feature of algorithms means they'll be poor models for a range of abilities associated with our own practical thinking. In particular, models designed to mimic the structure of computer algorithms won't neatly explain abilities that lack top-down structure. But are there such abilities? Are we able to thoughtfully do anything without having a well-defined aim in mind? Such questions have been explored in the literature on practical reasoning, so a quick glance there will help to address the question.

One standard vision of practical reasoning is instrumentalism,³² the view that desires are things we just have and that the only practical reasoning to be done given those desires is in determining how to fulfill them. Such a view is very much in line with the view of practical thinking articulated by Pavese. Just as computer "minds" execute algorithms to effect solutions for a given purpose, so, too, human minds operate by executing inferential

³⁰As an aside: even if you can specify the solution-space, delimiting the range of appropriate steps to achieve the solution may prove difficult. For example, you might reasonably believe that successfully babysitting a three year old means it will be unharmed when you arrive home. That solution, however, could be achieved by duct-taping the child to a chair for a few hours. Anticipating all the ways in which a solution can be achieved by adopting unacceptable means is a genuine problem, but not one I'll pursue here.

³¹Genetic algorithms may seem to lack this feature, since they are designed to produce solutions to problems through random mutation to a base (or set of base) algorithm(s). However, since mutation occurs (in the simplest and most general) case to bits of strings that are mapped to different problem domains, the problem-space the genetic algorithm is designed to address and what would count as a solution in that space must already be specified. Without a well-defined, clearly specified range of success conditions, determining which of a genetic algorithm's offspring meet those conditions would be impossible.

³²Instrumentalism is the default view and is often assumed rather than defended. R. Jay Wallace's (2014) discussion of practical reasoning in the Stanford Encyclopedia of Philosophy is a good starting place. Instrumentalism is typically traced to Hume, though there are good arguments for resisting that interpretation (see Millgram (1995)). Bernard Williams (1981) as well as Michael Smith (2004) articulate the basics of the position, and a more sophisticated version can be found in Frankfurt (1998). For an argument showing why instrumentalism is the default view, see Vogler (2002).

rules that satisfy given desires. For example, if I want to make an omelette, my ability to do so is explained by executing rules that will bring about the desire. But rules to effect ϕ can't be implemented if we don't yet know the nature of ϕ . If I want something unusual for breakfast, my desire is too vague, too underspecified to serve as a termination point in the execution of an algorithm. I must first specify the content of the desire—that is, I must first delimit a range of conditions that, if achieved, would count as a solution to the fulfillment of my desire—before it's possible to execute a series of rules to achieve it.³³

It's not unreasonable to believe that settling on a solution space often requires specification of this sort. Before we are able to do what we know how to do, we must first decide upon, or specify, the nature of what it is we want to do. This process of solution specification is itself something we are able to do, and, as such, needs to be explained by an account of practical thinking. What rules of inference, what algorithm, could explain our ability to specify or determine the solutions we hope to achieve? There are of course banal cases of deciding between an omelette or a pancake for breakfast, and in such cases, we might envision an algorithm or rule of inference that arbitrarily decides between two (or a range) of options. In such cases, a model of practical thinking designed to mimic computer algorithms might be okay, since it is only required to select among solutions that are already specified. But there isn't always (or even typically) a well-defined range of options to choose from. That is, deciding what we want to do or specifying a solution that would satisfy vaguely given hopes, desires, or ambitions isn't always a matter of clearing the facade from what is already there, as if what we really want is within a chunk of mental marble. Sometimes settling on a potential solution to a problem or specifying the nature of what we aim to do is a matter

³³For the difficulty involved in deciding on ends, see Millgram's (2008) discussion of the French ambition to design a world-class system of mass transportation. For a book-length treatment showing that figuring out our ends is a practical problem, see Millgram (1997). The problem of determining our ends is also taken up by Kolnai (1961-1962) and Smith (1994).

of genuinely innovating, inventing, or discovering something entirely new.³⁴

But if that is true, if specifying a solution space may require discovery, invention, or innovation, then the structural features of algorithms will make them a bad fit for modeling that process, since the cluster of abilities associated with creation, innovation, and discovery have the wrong structure. They lack sufficiently well-defined solutions, which is a prerequisite for designing an algorithm that could achieve them.³⁵ So, if we know how to specify, determine, innovate or invent well-defined solutions to genuinely new problems, and if our ability to do so demands explanation, we should look for a model of practical thinking that

³⁴Of course, it may be that this supposed ability is illusory. It *seems* as though human agents are able to innovate, invent, or discover things entirely new, but maybe that seeming is a mistake. If so, there's no obligation for Pavese's model (or any model for that matter) to provide an explanation. But to the extent that we have such abilities, a model like Pavese's owes something it won't be able to provide.

³⁵I want to stave off two associated worries. First, there is the temptation to go meta anytime a potential solution is indeterminate, underspecified, or vague. The idea is that settling the content of a solution space can be achieved by treating it as a component for the satisfaction of some higher-level (or more fundamental), well-specified aim. For example, if my aim is to have something unusual for breakfast, I may not be able to say specifically what that comes to and, consequently, I may be unable to follow a rule to achieve it. But I can specify the content of the higher-level aim to, say, satisfy my appetite. Since I know what would satisfy the content of that aim, and since I can follow a rule to achieve it, I can hold my higher-level aim fixed and specify the content of the lower by choosing among options that would fill it in as a poorly specified variable in a rule toward the achievement of a well-specified aim. The temptation to offer a solution to a lower-order problem by going meta, however, is typically an attempt to ignore a problem rather than solve it. You don't solve the problem of poorly specified ends by suggesting that at a high enough level, there must be an end that is well-defined. And besides, working up a hierarchy of aims to discover an ultimate, well-defined aim that could serve as the fixed point for filling in the content of all our other ends is a philosophical dream that has gone unsatisfied. (There are two ways of attempting to work up the hierarchy: first, locate a solution space that is sufficiently rich to ground all human pursuits as means to it; second, locate a solution space that is sufficiently rich to ground all my pursuits as those that are genuinely mine. The first avenue is what Aristotle and Mill were up to by suggesting that what everyone really wants (or should want) is one or another version of happiness. The second avenue is the Frankfurtian (Frankfurt, 1999, 1998; Bratman, 2007; Velleman, 2000) ideal of locating a most fundamental desire to serve as the foundation of agency.) It's also worth keeping in mind that what I've said in this footnote seems well-outside the philosophical aims of intellectualists.

Second, there is also the temptation to deny that discovery, innovation, or invention are things we really know how to do. On this line of objection, we can't design algorithms to discover, innovate, or invent new aims because doing so isn't really something agents have the ability to do. Rather, innovation, discovery, and invention are the product of luck, chance, or random guesses that just work out. Succumbing to this strategy for responding to the challenge of discovery, innovation, and invention to an account of practical thinking, however, is again an attempt to get around the problem by ignoring or denying it. It might turn out that we really don't need to explain the cluster of abilities associated with innovation, invention, and discovery, but if so, we'd need an argument showing why.

doesn't mimic the structure of computer algorithms. We should look for a model other than the one offered by intellectualists.

Finally, and as a passing worry for Pavese's model, let me note that there are important classes of abilities that are self-undermining when explained by rule-governed thinking. For example, my ability to make and keep friends, my ability to act spontaneously, or my ability to act lovingly toward my significant other are all undermined if I try to explain those abilities by appealing to rules. Indeed, if I try to explain my spontaneity by appealing to a set of rules by which I act spontaneously, you might get the impression that I'm delusional or just putting on an act. I'm not genuinely spontaneous if the things I do are the product of rules. And similarly for making and keeping friends or acting lovingly toward my spouse: if I use an algorithm to make and keep friends, or follow rules to act lovingly toward my significant other, the things I do will appear shallow, calculated, and desperate, which are features of action not compatible with stable relationships. Again, we should look for an alternative model of practical thinking if we want to explain such abilities.

4.4 Recap of the Problems

I have articulated three broad classes of abilities for which Pavese's model of practical thinking will be a bad explanatory fit. Each is designed to exploit slightly different presuppositions built into her model in order to reveal its limits for understanding the human mind. The things we are able to do often require thinking that doesn't nicely match the execution of algorithms: some of what we are able to do lacks the structure presupposed by algorithms; some of what we are able to do lacks the precision required of computer algorithms; and some of what we are able to do is done without a specifiable solution space in mind. In each case, the model of practical thinking at the heart of intellectualism will be out of place.

5 Conclusion

Intellectualists have put forward an array of arguments to defend the view that knowledge-how is a species of knowledge-that, but the plausibility of their view has always rested on the promise that there is a peculiarly practical way of entertaining information to produce action. For too long that promise was left unfulfilled. The recent arguments by Pavese are an impressive attempt to fill in the intellectualist vision of mind and discharge the obligation of that promise. But the strength of Pavese's view rests on the similarities between that vision and the mind of the agents for which she is trying to give an account. As I have argued, her view is thoroughly computational and its presuppositions infect her model of practical thinking. As a result, we're left with an account of the workings of our minds that is asymmetrical with a number of features of our own practical thinking. Consequently, her view is poorly positioned to explain many of our abilities.

I'm reminded of a passage from Hume's *Dialogues Concerning Natural Religion* (1948, 21-22). There he argues that the designs and machinations of men are far too removed from the order of the universe for us to reasonably draw any conclusions about it on their basis. The order we find in things of our own design and the order we find in the universe are too disparate to support the weight of inference by analogy. It seems to me similarly mistaken to draw inferences about the workings of the human mind by looking at the workings of one of its artifacts. So let me conclude by offering a couple of speculative comments for which I haven't argued directly.

The view of practical thinking articulated by Pavese rests on a view of the human mind that seems to me largely out of place. There is a dominant tendency in the most influential circles of philosophy to buy into a vision of the mind as a digital computer. That metaphor has provided and continues to provide tremendous value. But if our aim is to understand

ourselves, it is a perspective with severe limitations. The assumptions needed to model the mind in that way severely distort what it is that we're doing when we're thinking. That's not to say that such a model is without value or that it fails in every application. To the contrary. But if we are to avoid misunderstanding the workings of our own minds, we must remain cognizant of the fact that representations of it are mere models and that models involve distorting presuppositions. A computer might be an apt model for understanding certain features of our thinking, but using computers to arrive at a general understanding of the mind is tantamount to attempting to understand the general nature of an ancient culture by looking at its religious architecture. It may tell you something, but there will be much that gets missed. I hope I have given some support to these thoughts in this paper.

References

- Anscombe, G.E.M. 2000. *Intention*. Cambridge: Harvard University Press.
- Beck, Jacob. 2012. “The Generality Constraint and the Structure of Thought.” *Mind* 121:563–600.
- Bengson, John and Moffet, Marc. 2007. “Knowing-how and Concept Possession.” *Philosophical Studies* 136:31–57.
- . 2011. *Knowing How*. Oxford: Oxford University Press.
- Bratman, Michael. 2007. *Structures of Agency*. Oxford: Oxford University Press. ISBN 0195187709.
- Brookshear, J. Glenn. 2014. *Computer Science: an overview*. Boston: Pearson Education, 12 edition.
- Burge, Tyler. 2005. *Truth, Thought, Reason*. Oxford University Press.
- Danto, Arthur. 1965. “Basic Actions.” *American Philosophical Quarterly* 2.
- Davidson, Donald. 1980. “Actions, Reasons, and Causes.” In *Essays on Action and Events*, 3–19. Oxford: Clarendon Press.
- Downes, Steve. 1992. “The Importance of Models in Theorizing: A Deflationary Semantic View.” In D. Hull, M. Forbes, and K Okruhlik (eds.), *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, volume 1, 142–153. East Lansing, MI: Philosophy of Science Association.
- Dummett, Michael. 1993. *Frege: Philosophy of Language*. Cambridge: Harvard University Press, 2nd edition.
- Evans, Gareth. 1982. *Varieties of Reference*. Oxford: Clarendon Press.
- Fodor, Jerry. 1968. “The Appeal to Tacit Knowledge in Psychological Explanation.” *The Journal of Philosophy* 65:627–640.
- Ford, Anton, Hornsby, Jennifer, and Stoutland, Frederick (eds.). 2011. *Essays on Anscombe’s Intention*. Cambridge: Harvard University Press.
- Frankfurt, Harry. 1998. *The Importance of What We Care About*. Cambridge: Cambridge University Press.
- . 1999. *Necessity, Volition, and Love*. Cambridge University Press.
- Frege, Gottlob. 1948. “Sense and Reference.” *The Philosophical Review* 57:209–230.
- . 1956. “The Thought: A Logical Inquiry.” *Mind* 65:289–311.

- Hornsby, Jennifer. 2004. "Agency and Actions." In John Hyman and Helen Steward (eds.), *Agency and Action*, 1–23. Cambridge: Cambridge University Press.
- Hume, David. 1948. *Dialogues Concerning Natural Religion*. Hafner Library of Classics.
- Koethe, John. 2002. "Stanley and Williamson on Knowing How." *Journal of Philosophy* 99:325–328.
- Kolnai, Aurel. 1961–1962. "Deliberation Is of Ends." *Proceedings of the Aristotelian Society* 62.
- Korsgaard, Christine. 2008. *The Constitution of Agency*. Oxford: Oxford University Press.
- Lavin, Douglas. 2012. "Must There Be Basic Action." *Nous* 47:273–301.
- Millgram, Elijah. 1995. "Was Hume a Humean." *Hume Studies* 21:75–93.
- . 1997. *Practical Induction*. Cambridge: Harvard University Press.
- . 2008. "Specificationism." In J. Adler and L. Rips (eds.), *Reasoning: Studies of Human Inference and its Foundations*, 731–747. Cambridge University Press.
- . 2009. *Hard Truths*. West Sussex: Wiley-Blackwell.
- Morrison, Margaret. 1999. "Models as Autonomous Agents." In Mary Morgan and Margaret Morrison (eds.), *Models as Mediators*, 38–65. Cambridge: Cambridge University Press.
- O'Brien, Lucy. 2007. *Self-Knowing Agents*. Oxford University Press.
- Pavese, Carlotta. 2015a. "Knowing a Rule." *Philosophical Issues: A Supplement to Nous* 25:165–188.
- . 2015b. "Practical Senses." *Philosophers' Imprint* 15.
- . 2017. "A Theory of Practical Meaning." *Philosophical Topics* 45.
- Rödl, Sebastian. 2007. *Self-Consciousness*. Cambridge: Harvard University Press.
- . 2011. "Two Forms of Practical Knowledge and Their Unity." In Anton Ford, Jennifer Hornsby, and Frederick Stoutland (eds.), *Essays on Anscombe's Intention*, 211–241. Cambridge: Harvard University Press.
- Ryle, Gilbert. 1946. "Knowing How and Knowing That." *Proceedings of the Aristotelian Society* 46.
- . 1949. *The Concept of Mind*. New York: Barnes and Noble.
- Schaffer, Jonathan. 2007. "Knowing the Answer." *Philosophy and Phenomenological Research* 75:383–403.
- Schiffer, Stephen. 2002. "Amazing Knowledge." *Journal of Philosophy* 99:200–202.

- Schwenkler, John. 2015. "Understanding "Practical Knowledge"." *Philosophers' Imprint* 15.
- Searle, John. 1983. *Intentionality: an essay in the philosophy of mind*. Cambridge: Cambridge University Press.
- Setiya, Kieran. 2007. *Reasons without Rationalism*. Princeton: Princeton University Press.
- . 2008. "Practical Knowledge." *Ethics* 118:388–409.
- Shapiro, Stewart. 2008. *Vagueness in Context*. Oxford: Clarendon Press.
- Small, Will. 2012. "Practical Knowledge and the Structure of Action." In Günter Abel and James Conant (eds.), *Rethinking Epistemology*, volume 2, 133–228. Berlin: de Gruyter.
- Smith, Michael. 1994. *The Moral Problem*. Oxford: Blackwell.
- . 2004. "The Structure of Orthonomy." In John Hyman and Helen Steward (eds.), *Agency and Action*, 165–193. Cambridge: Cambridge University Press.
- Stanley, Jason. 2011a. *Know How*. Oxford: Oxford University Press.
- . 2011b. "Knowing (How)." *Nous* 45:207–238.
- Stanley, Jason and Williamson, Timothy. 2001. "Knowing How." *The Journal of Philosophy* 98:411–444.
- Suppes, Frederick. 1977. *The Structure of Scientific Theories*. Champaign: University of Illinois Press, 2nd edition.
- Tao Jiang, Ming Li, and Bala Ravikumar. 2004. "Formal Models and Computability." In Allen B. Tucker (ed.), *Computer Science Handbook*, chapter 6, 128–161. Chapman and Hall/CRC, 2nd edition.
- Thompson, Michael. 2008a. "Anscombe's *Intention* and Practical Knowledge." In Anton Ford, Jennifer Hornsby, Frederick Stoutland (ed.), *Essays on Anscombe's Intention*. Cambridge: Harvard University Press.
- . 2008b. *Life and Action*. Cambridge: Harvard University Press.
- Velleman, James David. 2000. *The Possibility of Practical Reason*. Oxford: Clarendon.
- Vogler, Candace A. 2002. *Reasonably Vicious*. Cambridge: Harvard University Press.
- Wallace, R. Jay. 2014. "Practical Reason." *The Stanford Encyclopedia of Philosophy* Edward N. Zalta (ed.).
- Weisberg, Michael. 2013. *Simulation and Similarity*. Oxford: Oxford University Press.
- Williams, Bernard. 1981. "Internal and External Reasons." In *Moral Luck*, 101–113. Cambridge: Cambridge University Press.

Williamson, Timothy. 1994. *Vagueness*. New York: Routledge.

—. 2002. “Epistemicist Models.” *Philosophy and Phenomenological Research* LXIV:143–150.

—. 2003. “Vagueness in Reality.” In Michael Loux and Dean Zimmerman (eds.), *The Oxford Handbook of Metaphysics*, 690–716. Oxford: Oxford University Press.