

Knowing What to Do*

Ethan Jerzak and Alexander W. Kocurek

Forthcoming in *Noûs*

Abstract

Much has been written on whether practical knowledge (knowledge-how) reduces to propositional knowledge (knowledge-that). Less attention has been paid to what we call **deliberative knowledge** (knowledge-to), i.e., knowledge ascriptions embedding other infinitival questions, like *where to meet*, *when to leave*, and *what to bring*. We offer an analysis of knowledge-to and argue on its basis that, regardless of whether knowledge-how reduces to knowledge-that, no such reduction of knowledge-to is forthcoming. Knowledge-to, unlike knowledge-that and knowledge-how, requires the agent to have formed certain conditional intentions. We discuss the philosophical implications for knowledge-how, deliberative questions, and virtue.

1 Introduction

Consider some things you can know: You can know that the stock market will crash soon, how to game a crashing market, what to do when it crashes, when to sell your stocks, who else to warn to do the same, and where to stash the gains once you do. Much attention in epistemology has been paid to the first of these kinds of knowledge, **propositional knowledge** (*knowledge-that*). There is also a vibrant debate on the second, **practical knowledge** (*knowledge-how*), especially concerning how it relates to propositional knowledge (e.g., Ryle 1949; Stanley and Williamson 2001; Noë 2005; Setiya 2008, 2012; Wiggins 2012; Abbott 2013; Pavese 2015, 2021; Santorio 2016; Habgood-Coote 2017; Constantin 2018; Boylan 2023). Much less philosophical work exists on the remainder. Our goal is to provide an analysis of them.

It is well-known that the verb *know* can take many types of clauses as arguments, including not just that-clauses but also interrogatives, noun phrases (both as concealed questions and as in acquaintance readings), and prepositional phrases:

*Thanks to Martín Abreu Zavaleta, Jens Kipper, Carlotta Pavese, Rachel Rudolph, Nic Southwood, and Zeynep Soysal for their feedback. Thanks also to the audiences of the Philosophy of Language Associate 2023, PhLiP 2023, the 2023 Hong Kong Expressivism Workshop, the Cornell Semantics Reading Group, and the Australian National University, where earlier versions of this material were presented. The authors contributed equally to this paper.

- (1)
 - a. Everyone knows that the earth is round.
 - b. Charlie knows who can sing falsetto.
 - c. Morgan knows the capital of Germany.
 - d. Luke knows [of] Patricia.

We are specifically interested in knowledge ascriptions that embed what's known as an **infinitival question** ("IQ" for short), i.e., an interrogative phrase embedding a full-infinitive verb phrase (*to* ϕ), such as *know who to call*, *when to meet*, *where to sit*, and *what to wear*. These contrast with knowledge ascriptions embedding a finite question like *know who is coming*, *when class is*, *where we are*, and *what time it is*.

Much of the literature on knowledge ascriptions embedding IQs has focused exclusively on *how to* questions—mostly on whether knowledge-how reduces to knowledge-that. Stanley and Williamson (2001) argue that it does: *S* knows how to ϕ iff there is a way w_ϕ (for *S*) to ϕ such that *S* knows that w_ϕ is a way (for *S*) to ϕ . Others have denied any such reduction is possible. Famously, Ryle (1949) argued that knowledge-how requires the possession of an ability, which cannot be reduced to knowing that such-and-such is the case. Others argue that knowledge-how is a non-propositional disposition of some kind, e.g., a disposition to be guided by one's intentions (Setiya, 2012) or to develop an ability (Constantin, 2018).

Other IQs (*what to*, *where to*, *when to*, etc.) have received only a passing mention in this debate (though see Farkas 2017 for a notable exception). Insofar as they are discussed, they are assumed to work much like *how to* questions: if knowledge-how reduces to knowledge-that, it is thought, then so does knowledge involving other IQs. Indeed, the standard analyses of *know*+IQ constructions in linguistics reduce them in some way to knowledge-that.

We think this is a mistake. Our aim is to show that, for reasons completely orthogonal to the issues concerning knowledge-how and abilities, these other knowledge ascriptions cannot be reduced to knowledge-that ascriptions.

Knowledge ascriptions embedding IQs tends to convey something about the decisions facing a subject and how that subject relates to them. When Riley knows what to wear to the ball, her knowledge concerns a sartorial choice she faces. When Kendall knows which wire to cut to defuse the bomb, his knowledge concerns a choice to cut the red wire or the blue wire. These are agents who *know what to do* when faced with the relevant choice. We will argue that the key to understanding *know*+IQ constructions (possibly excluding *how to*) is a more adequate analysis of *know to* constructions, which involve embedding a full-infinitive (without a *wh-* element) under *know*, as in (2):

- (2)
 - a. Riley knows to wear the blue dress.
 - b. Kendall knows to cut the red wire.

Agents know *what to do* in virtue of knowing *to do* something—they *know to* ϕ for some ϕ . For this reason, we'll refer to the kind of knowledge expressed with infinitival verb phrases (*to* ϕ) and IQs (*wh-* *to* ϕ , excluding *how to* ϕ) as **deliberative knowledge**, or *knowledge-to*.

Our argument against the reduction of knowledge-to to knowledge-that is basically this: knowledge-to entails that the agent has formed a certain kind of conditional intention. Specifically, it requires that the subject intend, of some course of action, to execute it when faced with a certain (possibly hypothetical) decision. This explains why agents who know what to do are both decisive and guided by their knowledge of what's best to do in the deliberative context. So long as these intentions themselves do not reduce to propositional knowledge (which, as we'll discuss, is contested but plausible), it follows that deliberative knowledge does not either.

We begin by recounting the existing analyses of deliberative knowledge from linguistics (section 2). We then present a problem for such analyses, and argue that this problem can be evaded only by appealing to controversial meta-normative principles (section 3). Thence follows our own analysis of deliberative knowledge in terms of knowledge-that plus conditional intentions (section 4). Some objections are entertained (section 5). We conclude by discussing some open questions for the study of deliberative knowledge as well as some of the philosophical implications for knowledge-how, deliberative questions, and the Socratic theory of virtue (section 6).

2 Infinitival Questions: A Primer

While knowledge-to has not received nearly as much attention as knowledge-how, it has received some. Largely, though, this attention has been paid in linguistics rather than philosophy. In this section, we'll briefly outline this literature in linguistics.

To start, let's consider knowledge ascriptions embedding finite questions. Consider (3):

(3) Libby knows where the library is.

On the standard analysis, (3) is true iff Libby knows the answer to the question *Where is the library?*. If the library is located at 135th St and Malcolm X Blvd, then Libby knows that the library is located there (perhaps under another guise or description). More generally, it is typically held that *S* knows *Q*, where *Q* is a question, iff *S* knows that *A*, where *A* is an answer to *Q*.¹ This fits with a more general analysis on which interrogatives denote a set of answers to the question expressed (Hamblin, 1958, 1973; Karttunen, 1977; Groenendijk and Stokhof, 1982; Ciardelli, Groenendijk, and Roelofsen, 2018).

Given this, it's natural to assume that knowledge ascriptions embedding IQs are no different: deliberative knowledge just amounts to knowing an answer

¹ There's some debate over whether and when the answer the subject knows must be complete and exhaustive (see, e.g., Braun 2006; Schaffer 2007; Mastro 2010; George 2013). This is complicated by the distinction between *mention-some* questions, whose answers need not be exhaustive (e.g., *Where can one buy an Italian newspaper?*), and *mention-all* questions, whose answers do need to be exhaustive (e.g., *Which students passed the exam?*). Here, we will use the term "answer" to remain neutral on this issue, as it will not be particularly central in what follows (though see section 5).

to the question expressed by the embedded IQ. But what *is* the answer to an IQ? We do not typically go around asking questions stated directly with such infinitival clauses, like *Which dessert to eat?* or *Where to buy a phone?*. We might mutter such questions to ourselves, but they would be very odd to pose to someone else. When directed at others, we are more likely to state the question without an infinitive, like *Which dessert should I eat?* or *Where can one buy a phone?*. How, then, should we analyze IQs and their “answers”?

The standard analysis of IQs in linguistics is arguably due to Bhatt (1999). Bhatt argues that IQs can be analyzed with a covert modal, which can be paraphrased as either *should*, *could*, or (less commonly) *would* depending on a number of factors. So we can paraphrase (4)–(5), for example, roughly along the following lines:

- (4) Magnus knows *where to get gas*.
 \approx Magnus knows *where he/one can get gas*.
- (5) Hafdis knows *who to talk to*.
 \approx Hafdis knows *who she/one should talk to*.

Bhatt introduces a new practical (deontic) modal, $\diamond_{D,\rightarrow}$, which is used as the covert modal in IQs. Roughly, $\diamond_{D,\rightarrow} p$ says that p can satisfy the contextually relevant goals and does so at every contextually relevant world. Thus, (4) is analyzed along the following lines:

- (6) Magnus knows [where _{i} $\diamond_{D,\rightarrow}$ (PRO to get gas t_i)]
 \approx Magnus knows that he/one can get gas at t_i and that getting gas at t_i suffices to achieve the contextually relevant goals.

Here, the subject of the IQ is an unpronounced PRO, which can either be interpreted in a de se way (where *Magnus* can get gas) or in a generic way (where *one* can get gas). The trace element t_i is indicating syntactic movement. So on the de se reading, for example, the denotation of *where to get gas* is the set of propositions of the form *Magnus can get gas at t_i* . Magnus knows where to get gas iff he knows one of the propositions in this set.

Bhatt suggests that the force of this modal $\diamond_{D,\rightarrow}$, e.g., whether it’s best paraphrased as *should* or *could/can*, depends on a variety of factors. For example, the *wh*-element can affect the interpretation of this modal. So while *where* and *how* often are paraphraseable as *could/can*, it’s more natural to paraphrase *who*, *when*, *whether*, and *which* using *should*.

- (7) a. Burt knows where to find a good burrito.
 \approx Burt knows where he can find a good burrito.
 b. Rose knows how to solve the equation.
 \approx Rose knows how she can solve the equation.
- (8) a. Jackie knows who to invite.
 \approx Jackie knows who she should invite.
 b. Maya knows when to interrupt.
 \approx Maya knows when she should interrupt.

- c. Trystan knows whether to turn left or right.
 \approx Trystan knows whether he should turn left or right.
- d. Fiona knows which route to take.
 \approx Fiona knows which route she should take.

These are just general tendencies, not ironclad rules. Context plays a large role in determining the interpretation of an IQ. If it's understood that multiple options can and would equally satisfy all the relevant goals, then the *could* reading will be more salient. If, however, it's assumed that there will only be one unique option that satisfies all the relevant goals, then the *should* reading will be preferred.

The exact details of when IQs receive which reading are not particularly important for our purposes. What's more important is that on Bhatt's account, deliberative knowledge reduces to propositional knowledge: *S* knows what to do iff *S* knows what *S* can and would (at the contextually relevant worlds) satisfy the (contextually relevant) goals, i.e., there is a ϕ such that *S* knows that *S* can/should ϕ given said goals.²

An alternative analysis of IQs was put forward by Roberts (2009). Roberts observes that knowledge-how ascriptions have what she calls a certain "de se character". Imagine you're looking at a map of hiking trails but you do not know where on the map you are (it lacks a "You are here" label). Even though you know how one could get back to the welcome center from any location on the map, you do not know how to get back to the welcome center yourself.

For reasons that need not detain us, Roberts argues that this essentially de se character of knowledge-how ascriptions (and *know*+IQ more generally) is better captured by analyzing IQs as properties, rather than as propositions with a hidden PRO subject. Following Dowty and Jacobson (1991), Roberts analyzes IQs as effectively the predicate-analogues of interrogatives: IQs are to other interrogatives as predicates are to declarative sentences. On this view, IQs denote functions from agents to questions,³ specifically questions concerning what goals it would be rational for the agent to adopt under ideal circumstances.⁴ For example, she analyzes (9) as follows:

- (9) James knows what to drink with spicy food.
 \approx James knows the *y* such that *drink y with spicy food* would be a rational goal to pursue in ideal circumstances.

² The 'can' here is deontic (as in 'is permitted to [given the relevant goals]'), not circumstantial (as in 'is able to'). So if Magnus doesn't want to commit a crime and only knows where one can get gas illegally, he doesn't know where to get gas. However, Bhatt suggests $\diamond_{D,\rightarrow} p$ can reduce to a circumstantial claim if it's "trivially true" that every relevant *p*-world satisfies the relevant goals.

³ Roberts follows the standard analysis of questions, due to Groenendijk and Stokhof (1982), as partitions on worlds, where each cell of the partition agrees on the answer to the question (cf. Lewis 1988a,b).

⁴ By "ideal circumstances", Roberts has in mind circumstances in which there is a "lack of conflict with other, pre-established goals, commitments and intentions of the agent". By "rational goal", she means a goal that (i) potentially will have a payoff, (ii) one has reason to think is achievable, and (iii) doesn't conflict with the agent's other goals or commitments.

The exact details of how to flesh this out are not important for our purposes. What's important is that, like Bhatt's account, Roberts's account reduces deliberative knowledge to propositional knowledge: *S* knows what to do iff for some ϕ , *S* knows that it would be rational to ϕ (in ideal circumstances).

3 The Deliberative-Normative Divide

Both of these accounts analyze IQs as disguised normative questions concerning what can or should be done to satisfy certain goals. For Bhatt, knowing what to do is the same as knowing what one *can/should* do in order to satisfy the contextually relevant goals. For Roberts, it's knowing what goal would be *rational* to pursue in ideal circumstances.

But is that right? Is the question of what to do the same as the question of what one should do? According to a prominent strand of the literature in metaethics, the answer is no: the **deliberative question** of *what to do* differs from the **normative question** of *what one can/should do*, even if we interpret *can* and *should* prudentially (Bratman, 1987; Moran, 2001; Broome, 2008; Hieronymi, 2009, 2011; Owens, 2011; Southwood, 2016; Balcerak Jackson, 2019; Clarke-Doane, 2020, 2022; Risberg, 2023).⁵

One potential difference between deliberative and normative questions concerns how they are resolved. To resolve a normative question, one only needs to form a *belief* of a certain kind: a belief that one can or should do such-and-such. To resolve a deliberative question, however, one must form an *intention* or make a *decision*. Yet forming a normative belief does not necessarily involve forming the requisite intention or vice versa. There are two kinds of cases that are typically used to illustrate this: cases of **akrasia** (cf. weakness of will) and cases of **parity** (e.g., Buridan's ass).⁶

In cases of akrasia, an agent decides to do something they think they shouldn't do. For example, maybe Akira knows she should study for her exam tomorrow but decides to play *Call of Duty* instead. Thus, Akira has answered the normative question of what she should do one way (study), but answered the deliberative question of what to do a different way (play video games).

In cases of parity, an agent considers multiple options to be equally good (or "on par"), but still must decide between them. For example, maybe Parry knows that he should go to either Ann's party or Ben's party, and that both are permissible, but hasn't decide which to go to. Thus, Parry has answered all the relevant normative questions (he should go to one of the parties; it's not the case he should go to Ann's and also not the case he should go to Ben's), but has yet to answer the deliberative question of which party to go to.

⁵ Deliberative questions are sometimes called "practical questions", especially in the literature on "practical" reasoning.

⁶ Two other cases (similar to parity) could also be used here. First, there's the "consistent normative nihilist", who believes there is nothing they (or anyone) should do (Southwood 2016; see Faraci 2017). Second, there are moral dilemmas where an agent is normatively required to do incompatible things (thanks to Al Hájek for suggesting this). In both cases, agents can still ask and consider deliberative questions even though they've answered the relevant normative questions.

The possibility of these cases is contested. Critics argue these cases are misdescribed (Hare, 1952; Davidson, 1980; Gibbard, 2003, 2008; Chislenko, 2016). Akira doesn't *really* believe she should study all things considered if she decides to play video games. And while Parry might believe *before* he makes a decision that the two options are equally good, *after* he decides to go to Ann's party (say), he no longer views going to Ben's party as equally good. Advocates reply these cases are not misdescribed (Bratman 1987; Broome 2008; Hieronymi 2009; cf. Ullmann-Margalit and Morgenbesser 1977; Chang 2002). Akira can genuinely believe (and even know!) that she should study while deciding not to. And Parry could choose to go to Ann's party while genuinely viewing going to Ben's party as an equally good alternative.

Our goal is not to settle this long-standing philosophical debate here. Instead, we want to use this debate to formulate a desideratum on an adequate theory of deliberative knowledge. Arguably, a semantic analysis of deliberative knowledge attributions should not rest on a controversial metanormative principle relating deliberative and normative questions. So we should prefer an analysis that does not stand or fall with the possibility of akrasia or parity. That means an analysis should at least be *compatible* with views according to which the practical and the normative come apart, in that they yield the right results in cases like akrasia and parity where the deliberative and normative diverge.

We propose, then, that an adequate analysis of deliberative knowledge meet the following requirement:

Desideratum: Philosophical Neutrality

An analysis of *know*+IQ should yield the right results in cases of akrasia and parity, assuming they're possible.

Yet, as we will now argue, it is precisely these sorts of cases on which Bhatt's and Roberts' analyses founder.

Consider Akira before she decides to play *Call of Duty*. She knows that she should study for her exams. And yet, she is tempted to play video games instead. In this scenario, consider (10):

(10) Akira knows what to do.

In some important sense, (10) seems false. After all, Akira hasn't made up her mind yet! We can even imagine Akira saying to herself, "*I really* should study, but I'm torn. . . Ughhh, I don't know what to do!". She hasn't figured out what to do, and so doesn't know what to do. Yet both Bhatt's and Roberts's accounts predict that (10) is unequivocally true.

To break this down, consider Bhatt's account first, which says (10) is true iff there is something Akira can do that she knows can and would satisfy the contextually relevant goals (at the contextually relevant worlds). The righthand side is satisfied: Akira knows that studying can and would satisfy her goals of, say, passing her exams, getting good grades, etc. In other words, she knows she should study given her goals. So according to Bhatt's account, (10) is true.

Now consider Roberts's account. On this account, (10) is true iff there is something Akira can do that she knows would be a rational goal to pursue in ideal circumstances. Again, the righthand side is satisfied: Akira knows that it would be rational for her to study in ideal circumstances (of course, she's not in such circumstances, but that doesn't matter: she only needs to know what would be rational ideally). So according to Roberts, (10) is true.

One might object that Akira does not *really* know that she should study, or that it would be rational to do so: her behavior belies her claim to normative knowledge. But again, we are working under the assumption that akrasia is possible: we want an analysis of deliberative knowledge that accommodates these cases. If one does not consider this a genuine example of akrasia, one can swap in their preferred example. As long as any case with this structure is genuinely possible, Bhatt and Roberts are in trouble.

One might instead object that Akira only *believes* (correctly) that she should study, or that it would be rational to do so, but her true belief doesn't constitute knowledge. We have a hard time seeing why, though. We can imagine Akira has the strongest evidence for her belief: she has always passed her exams when she studied and not otherwise, she doesn't know the material that well yet, etc. Further, her akratic intentions do not necessarily constitute a defeater: she may know she's acted akratically in the past, intending to do what she shouldn't do, and have good reason for thinking she's about to do it again. In principle, then, an akratic agent ought to be able to not just *correctly believe* but also *know* they should ϕ , even while intending not to do so.

Finally, one might insist that in this case, it isn't even *true* that she should study or would be rational to do so. After all, Akira seems to have multiple goals that conflict, e.g., to pass her exam *and* to have fun. Nothing she does can satisfy all the contextually relevant goals: she's just in a normative dilemma. The trouble with this response, however, is that it overcorrects. Imagine that after struggling with her decision, Akira decides to study after all. In that case, (10) seems unequivocally true. Yet if there's *nothing* Akira can or should do to given her goals, if *nothing* would be rational for her to do, then according to both Bhatt and Roberts, (10) is automatically false.

Now, in fairness, there does seem to be another sense in which (10) is true. It does not sound totally contradictory to say "Akira knows what to do. She just can't/won't do it." Later, we'll consider some ways of explaining these two readings (section 5). It is possible Bhatt's and Roberts's accounts can adequately capture this alternate reading (though we will raise some independent doubts). For our purposes, it suffices that there's at least *a* reading on which (10) is false. In other words, there's at least one reading, which their accounts do not capture, on which an indecisive agent does not know what to do.

Parity poses further problems. Recall Parry, who is trying to decide whether to go to Ann's or Ben's party. He knows that he should go to one, that he is able and permitted to go to either, and that he can't go to both. He finds himself struggling to decide. In this scenario, consider (11):

(11) Parry knows what to do.

In some sense, (11) is true: Parry knows what to do, viz., go to one of the parties. But in another sense, (11) is false: he doesn't know whether to go to Ann's party or Ben's party. Yet, once again, both Bhatt's and Roberts's accounts predict (11) is true full stop. Parry knows what can and would satisfy his goals, what goals would be rational to adopt, and so on.⁷

As before, one might object that Parry doesn't *really* know what can and would satisfy his goals or be rational to do. But again, we're assuming that parity cases are possible. Or one might object that Parry believes but doesn't know that, say, going to Ann's party would be rational. But again, we fail to see why, given his evidence is sufficiently strong. Or perhaps one might object that nothing can satisfy Parry's conflicting goals. But again, this overcorrects, predicting (11) is false once Parry decides to go to (say) Ann's party.

Both of these cases suggest that deliberative knowledge requires more than knowing what one can or should do, or what would be rational to do. Such knowledge is theoretical and doesn't require an agent to act on it. To know what to do, one must *decide* to do one of the things that they should, or that it would be rational to, do.

4 Analysis

We now present our own positive analysis of deliberative knowledge. Let's return to an idea stated in section 2, namely that knowing a question amounts to knowing an answer to it. If we can analyze what it is to know an answer to a deliberative question, we can bootstrap this to an analysis of deliberative knowledge more generally.

For Bhatt and Roberts, the answer to a deliberative question (for an agent) is a proposition concerning what it is rational for the agent to do, what they should or could do, etc. But akrasia and parity show that something else is required to answer a deliberative question; agents can answer the question what they should do, without thereby answering the question what to do. To answer such a question, one needs to do more than just know that one should do something; one must actually decide or intend to do it. This suggests that knowing the answer to an IQ is not merely a matter of *knowing that* some proposition is true. But then, what *does* knowing the answer to an IQ involve?

We suggest it involves *knowing to* ϕ for some ϕ . There has been almost no discussion of *know-to* constructions in the literature (though see brief remarks by Wiggins (2012, p. 113)). But they are fairly natural in ordinary language:

⁷ Bhatt (1999, pp. 156–158) seems aware of this problem. He suggests that in these cases, $\diamond_{D,\rightarrow}$ is more naturally paraphrased as *should*, though it's not clear how his semantics can predict this. (He explains how *S knows what they should do* can be true given *S knows they should either ϕ_1 or ... or ϕ_n* is true, but never shows how this pertains to *S knows what to do*; he simply presupposes *S knows what to do* is equivalent to *S knows what they should do* here.) His semantic clause for $\diamond_{D,\rightarrow}$ does not seem to license the *should*-paraphrase in these cases. Moreover, it's not clear Bhatt can consistently maintain it does. Earlier (pp. 152–154), he says the *should*-paraphrase primarily arises when there's exactly one option that can and would satisfy the relevant goals. By definition, however, parity cases are precisely those where this condition is not met.

- (12) Reed knows to bring the reading to class.
- (13) Fiora knows to take the stairs in the event there's a fire in the building.
- (14) Chris knows to look both ways when crossing the street.

We think that these constructions must form the backbone of any analysis of deliberative knowledge, for the knowledge they attribute is precisely the knowledge of the answers to the deliberative questions posed by IQs.

Our strategy therefore consists of three stages. Stage 1 will develop an analysis of *S knows to ϕ* .⁸ Stage 2 puts forward an analysis of *S knows what to do* in terms of knowing to ϕ . Stage 3 extends this analysis to *know+IQ* constructions generally.

Stage 1: *S knows to ϕ* . While neither Bhatt nor Roberts discuss *know to* constructions directly, we can infer a natural extension of their analysis to such constructions. Arguably, Bhatt would say that *S knows to ϕ* iff *S* knows that ϕ -ing can and would satisfy the contextually relevant goals. And Roberts would say that *S knows to ϕ* iff *S* knows that ϕ -ing would be a rational goal to pursue in ideal circumstances. Note that both accounts reduce these knowledge ascriptions to propositional knowledge.

We agree that *know to* constructions require some propositional knowledge. Specifically, we think that all agents who know what to do know what they should do.⁹

But we claim that *know to* constructions also require something further: they also require that the subject form a certain kind of intention. Not only do they need to know what they should do in the face of a certain decision, they also need to intend to do it when faced with that decision.¹⁰ This is why, for example, Akira and Parry lack deliberative knowledge in their respective cases: even when they know the normative facts, they still need to make a decision.

To be clear, our claim is not that knowing to ϕ implies intending to ϕ full stop. One can know to ϕ when ϕ -ing is not even an option. Imagine Reed learns his textbook was stolen and he has no other means of acquiring another copy before class. In that case, (12) still sounds true—Reed knows to bring the

⁸ There are some *know to* constructions we wish to set aside. First, there are “de re” cases where the infinitive results from movement outside the scope of *know* (*Ben said something he knows to be false*). Second, there are cases where the infinitive modifies the prejacent, which are usually paraphrased with *in order to* (*Ann knows [in order] to get an A she must study*). None of these cases strike us as distinctively deliberative and arguably can be analyzed in terms of knowledge-that.

⁹ We leave open what flavor of ‘should’ (prudential, moral, etc.) is used here. While both Bhatt and Roberts seem to invoke prudential ‘should’ and ‘can’, we think that some *know to* constructions could be understood morally in cases where the two come apart. (Of course, Bhatt and Roberts could accommodate such cases by simply specifying the relevant goals to include moral goals.)

¹⁰ Farkas (2017) suggests a related account of “practical knowledge-wh” (including what we call deliberative knowledge), which requires the agent to be able to “activate knowledge of an answer in an appropriate context where the question arises” (p. 867). This would be insufficient to explain the kinds of cases discussed in section 3, however, since agents can possess such an ability while remaining indecisive.

reading to class—even if Reed doesn’t intend to bring the reading to class since he knows he can’t. However, *if* Reed can and should bring his textbook to class, then (12) requires he intend to do so.

To capture this, we assume agents adopt intentions relative to a (possibly hypothetical) task or decision situation (cf. Cariani, Kaufmann, and Kaufmann 2013 on deontic modals). Knowledge-to ascriptions often mention such decision situations explicitly, as illustrated in (13) and (14) (repeated below):

(15) Fiora knows to take the stairs in the event there’s a fire in the building.

(16) Chris knows to look both ways when crossing the street.

The decision-relativity of knowledge-to ascriptions can help explain how Reed can still know to bring his textbook even once he’s learned it was stolen. While no decision is explicitly mentioned, it is arguably contextually understood that the decision problem pertaining to (12) is the decision of whether to bring the textbook, where that is a genuine option for Reed. In other words, what Reed knows is to bring the textbook *when faced with a certain decision* where bringing the textbook is a genuine option. It’s just that, unfortunately, the decision is not one Reed faces since his textbook was stolen. He only has deliberative knowledge relative to an entirely *hypothetical* decision situation.

We propose, then, that all *know to* constructions make reference (either explicitly or via context) to a decision situation—that is, the construction we aim to analyze is *S knows to ϕ in D* , where D denotes a decision situation. Formally, we can think of a decision situation for an agent S as a set $\{\phi_1, \dots, \phi_n\}$ of options S can pick from (cf. Weirich 1983; Hedden 2012; Schwarz 2021).¹¹

Our claim is that in order for S to know to ϕ in a decision situation D , they must satisfy two conditions. First, S must intend to [ϕ in D]. This is a kind of conditional intention: it is the intention to ϕ *given* one is faced with decision D . Second, S must also know they should [ϕ in D]. This is a kind of conditional deontic proposition: one ought to ϕ *given* one is faced with decision D .

Here, then, is our analysis:

Analysis of *know to*

S knows to ϕ in D iff (1) S knows they should ϕ in D , and (2) S intends to ϕ in D .

Where no decision situation is specified, as in (12), we assume context supplies one. Usually (though not always), this is the decision situation that S faces at the time of evaluation.

For example, suppose Reed currently has two options: bring the reading to class or leave the reading behind. Then we can analyze (12):

¹¹ Here, we set aside a number of philosophical questions concerning the nature of options. For example, we remain neutral on whether something can count as an available option in a decision situation when the agent is unaware of it or doesn’t know how to do it (cf. section 6). We only assume that they are the kinds of things an agent can decide to do.

- (17) Reed knows to bring the reading to class.
 \approx Reed knows he should bring the reading to class in his current situation and he intends to bring the reading to class.

If Reed's textbook is stolen, however, then we should replace *his current situation* in the paraphrase above with a description of a hypothetical decision where bringing the reading is a genuine option for Reed. Similarly, we can analyze (13) as follows:¹²

- (18) Fiora knows to take the stairs in the event there's a fire in the building.
 \approx Fiora knows she should take the stairs in the event of a building fire and Fiora intends to take the stairs in such an event.

Stage 2: *S* knows what to do. To bootstrap this analysis of *know to* into an analysis of *know+IQ*, we start with two key ideas. First, *know+IQ* exhibits the same decision-relativity as *know to*. So our target is to analyze the construction: *S* knows *wh-* to ϕ in *D*. To work up to our analysis, we'll first state our analysis for *know what to do*, which avoids certain complications required for other IQs.

Second, knowing a question can be analyzed in terms of knowing an answer. To know an answer to a question, one must at least "accept" that answer. For finite questions, this simply amounts to believing a proposition. For IQs, however, more is required (if we assume deliberative and normative questions come apart): to accept an answer to an IQ, one must decide to do something in the relevant decision situation. As Southwood (2016) puts it, "we resolve the question of what we ought to do by forming a normative *belief*, whereas we resolve the question of what to do by forming an *intention* or *decision*".

The answer to an IQ, then, is not a proposition but a set of options (cf. Dowty and Jacobson 1991; Roberts 2009). So the IQ *what to do in D* denotes the set of options that are compatible with what the agent should do in *D*. For example, if Parry's options are {go to Ann's party, go to Ben's party, stay home}, where he should go to one of the parties but is permitted to go to either, then *what to do in D* denotes the set {go to Ann's party, go to Ben's party}.

As a first pass, then, one might propose the following:

Analysis of *know what to do* (first pass)

S knows what to do in *D* iff for some option $\phi \in D$, *S* knows to ϕ in *D*.

This account yields the correct results in many deliberative knowledge ascriptions. Fiora knows what to do in the event of a fire iff there's something she can decide to do in such an event (e.g., take the stairs) such that she knows to do that. It also makes the correct prediction about akrasia cases: Akira does not know what to do since Akira does not intend to study.

¹² Here, it may be more appropriate to think of *the event there's a building fire* as denoting not a single specific decision situation but a generic decision situation, or perhaps a class of decision situations, which may vary on which options exactly are available to Fiora, but all contain the option to take the stairs.

However, this account is inadequate for parity cases, when there are multiple options compatible with what the agent should do. Suppose Parry knows he could just as well go to Ann's party or Ben's party, but arbitrarily decides to go to Ann's. In that case, Parry arguably knows what to do in every relevant sense: it would be incorrect describe Parry as not knowing what to do once he's made a decision. But then the first pass account predicts the following is true:

(19) Parry knows to go to Ann's party.

This seems incorrect. One way to see this is that given just a little logical reasoning on Parry's part, (19) implies:

(20) Parry knows not to go to Ben's party.

Yet the conclusion sounds plainly false: going to Ben's is a perfectly fine option! Because going to Ann's party is not the *only* thing compatible with what he should do, Parry can't know to go to Ann's party and not to Ben's. In other words, in cases of parity, someone can know what to do without knowing to do some specific option, contrary to the first pass account.

As a second pass, one might propose the following:

Analysis of *know what to do* (second pass)

S knows what to do in D iff S knows to ϕ in D , where ϕ is the disjunction of options ψ_1, \dots, ψ_n that are compatible with what S should do in D .

Thus, after deciding to go to Ann's party, Parry counts as knowing what to do since he knows he should go to either Ann's or Ben's party and he intends to do just that. Note that if $n = 1$, then this reduces to the first-pass analysis.

Unfortunately, this account overcorrects: it predicts that Parry knows what to do even *before* he decides to go to Ann's party. And while that does seem right in some sense, there still seems to be another sense in which Parry does not know what to do precisely because he hasn't yet decided which party to go to: he doesn't know whether to go to Ann's party or Ben's.

The first pass account gets the right result for Parry before he makes his decision, but not after. The second pass account gets the right result after his decision, but not before. Intuitively, then, we need some "mixture" of these two accounts. The reason Parry knows what to do once he's made the decision is that he knows what he should do (go to one of the parties) and, moreover, he intends to do something that's compatible with that (go to Ann's party).

This suggests the following account, which is the one we endorse:

Analysis of *know what to do* (final)

S knows what to do in D iff (1) S knows to ϕ in D , where ϕ is the disjunction of the options ψ_1, \dots, ψ_n that are compatible with what S should do in D , and (2) for some ψ_i , S intends to ψ_i in D .

On this account, knowing a question still amounts to knowing an answer to that question. It simply generalizes this idea to IQs, where "answering" or "resolving" the question requires forming an intention, rather than simply

forming a belief. Thus, Parry does not answer the question of what to do, given the choice between {go to Ann's party, go to Ben's party, stay home} until he has made a decision as to which of those to do. Once he resolves to go to Ann's party, he knows what to do, since (1) he knows he should go to either Ann's party or Ben's party, which is all that's required of him in this scenario, and (2) he intends to go to Ann's party.

As desired, this account predicts that *S knows what to do in D* does not necessarily imply *S knows to ϕ in D* for some option ϕ in *D*. Even after Parry decides to go to Ann's party, he does not thereby know to go to Ann's party. Parry knows to go to Ann's party only if he knows he should; yet while he should go to Ann's party or Ben's party, neither is required. Even so, once Parry decides, he knows what to do: he knows what he should do and he has decided to do something compatible with that.

Our account does predict that *S knows what to do in D* does imply some statement of the form *S knows to ϕ in D*. It's just that the ϕ in question may not be a specific option in *D*: it could be a disjunction of options. Thus, independent of whether Parry decides which party to attend, the following inference is licensed:

- (21) Parry knows what to do.
⇒ Parry knows to go to either Ann's party or Ben's party.

Thus, on our account, someone knows what to do in virtue of knowing to do something or other. But how specific that "something or other" must be depends on how finely we individuate the available options (see section 5).

Stage 3: *S knows IQ*. What about other IQs, like *what to wear, when to leave, where to meet*, and so on? Here, we loosely follow Bhatt (1999, p.143ff) and propose that more specific IQs restrict the options under consideration in the decision situation. To illustrate, consider the following:

- (22) Sandy knows where to get a good sandwich during the lunch break.

The phrase *during the lunch break* suggests the decision situation is broadly a decision of what to get for lunch, which may include options other than grabbing a sandwich. Even so, (22) only seems to concern the sandwich-options: it only requires Sandy to know which sandwich options one should pick from *if* she decides to get a sandwich. It does not imply Sandy knows she should get a sandwich period. That is, (22) does not imply:

- (23) Sandy knows to get a sandwich during the lunch break.

However, if the deli is the best sandwich option in town, then (22) does imply the following:

- (24) Sandy knows to get a sandwich at the deli if she were going to get a sandwich during the lunch break.

What this suggests is that even if the decision situation in question includes a broad range of decisions, the IQ restricts the choice to just the options of a certain form. For example, *where to ϕ* restricts the options in the relevant decision situation to those of the form *ϕ in location l* (or perhaps *ϕ in locations l_1, \dots, l_n*). Similarly, *when to ϕ* restricts the options to those of the form *ϕ at time t* (or *ϕ at times t_1, \dots, t_n*). And so on.

To make this more precise, where Q is an IQ (not involving *how*; see section 6), let's write D_Q for the set of options in D that are of the "appropriate form". Specifically, for each Q below, D_Q is the set of options in D that are of the following form (or are positive boolean combinations thereof):

<i>who/what to ϕ</i> :	ϕx , where x is an object (or group of objects) of the appropriate type
<i>which/what Fs to ϕ</i> :	ϕx , where x is an F (or group of F s)
<i>where to ϕ</i> :	ϕ in l , where l is a location
<i>when to ϕ</i> :	ϕ at t , where t is a time (or interval)
<i>why to ϕ</i> :	ϕ for r , where r is a reason to ϕ
<i>whether to ϕ or ψ</i> :	ϕ, ψ
<i>how F to ϕ</i> :	ϕ to d , where d is a degree of F to ϕ

We allow D_Q to consist of options in D that are *positive boolean combinations* (e.g., disjunctions, conjunctions, disjunctions of conjunctions, etc.) of options of the corresponding form since the decision situation D may include some coarse-grained disjunctive options like "get gas at one of the following locations. . ." or "leave for work at one of the following times. . .", or conjunctive options like "call Callie and Callum" or "invest in each of these stocks". Such options seem to be left open by the corresponding *wh*-element. For example, Janet may know when to leave for work, where the options are leave between 8–9am, between 9–10am, and so on, in virtue of knowing to leave between 8–9am. Similarly, Magnus may know where to get gas, where the options are to get gas at one of these stations or at one of those stations, etc., in virtue of knowing to get gas at ones of those stations.

Note we include *how F to ϕ* constructions, as in *how many books to read*, *how much milk to buy*, *how friendly to be*, and *how high to climb*. These seem more naturally interpreted as deliberative knowledge rather than as practical knowledge. For example, *how many books to read* restricts a decision situation to those options of the form *read n -many books*, *how much milk to buy* restricts it to those options of the form *buy x -much milk*, *how friendly to be* restricts it to those options of the form *be friendly to degree f* , and so on.

Our more general analysis of *know*+IQ can be stated as follows:

Analysis of *know*+IQ

S knows IQ in D iff (1) S knows to ϕ in D_{IQ} , where ϕ is the disjunction of the options ψ_1, \dots, ψ_n that are compatible with what S should do in D_{IQ} , and (2) for some ψ_i , S intends to ψ_i in D_{IQ} .

This account nicely generalizes our account of *knows what to do*, since *what to do* carries a trivial restriction to the (disjunctions of) options of the form *do ϕ*

where ϕ is an action (or a set of actions). By contrast, *what to wear* carries a non-trivial restriction to the (disjunctions of) options of the form *wear x* where x is a wearable item (or a set thereof).

As predicted, *know+IQ* does not entail the corresponding *know to* claim. Thus, Sandy knows where to get a good sandwich during the lunch break iff Sandy knows which locations she should get a sandwich from *given* she's going to get a sandwich, i.e., relative to the choice of *where to get a sandwich* during the lunch break, and she intends to get a sandwich at one of those locations *given* she's going to get a sandwich. This is not the same as saying Sandy knows to get a sandwich during the lunch break, which requires her to know that she should (and that she intend to) get a sandwich period.

5 Discussion

In this section, we discuss some clarifications of and potential objections to our analysis of deliberative knowledge.

Akrasia Revisited. In section 3, we argued that in cases of akrasia, both Bhatt's and Roberts's accounts incorrectly predict that the akratic agent knows what to do. However, there do appear to be readings on which the following sound true in Akira's situation:

- (25) Akira knows what to do. . .
 - a. . . she just won't do it!
 - b. . . she just can't bring herself to do it!

We agree that there's at least *a* reading of *know what to do* where this is correct. We do not think this undermines our analysis, however.

First, even if akrasia does not pose a problem for extant analyses of *know+IQ* constructions, parity still does. Parry knows what he should do, what he can do, and what it would be rational to do, but still doesn't know what *to* do precisely because he has yet to make a decision. So parity cases already motivate the need for an analysis like the one we've provided.

Second, even if there is a sense in which Akira knows what to do (but can't/won't do it), there does seem to be a clear sense in which Akira does not know what to do (because she's indecisive). Again, we can imagine Akira muttering to herself: "I know studying would be the best thing, but I *really* would rather play *Call of Duty*. . . ugh, I don't know what to do!" And it is natural to describe Akira by saying, "Akira knows she needs to study to pass her exams. Yet, she's clearly waffling: she doesn't know whether to play video games instead." There is a prominent reading of knowledge-to ascriptions on which they're incompatible with indecision. This is a reading that the current literature ignores, and which, we've argued, cannot reduce to propositional knowledge.

There are several ways to explain the presence of multiple readings here. One possible explanation returns to the distinction between *de se* readings and *generic* readings of IQs. Recall that there seem to be two readings of the following:

- (26) Magnus knows where to get gas.
- a. \approx Magnus knows where *he* can get gas. (*de se*)
 - b. \approx Magnus knows where *one* can get gas. (generic)

This *de se*/generic distinction is reflected in knowledge-how ascriptions as well. Suppose Brooke has passed her driving test with flying colors, but breaks her leg, leaving her unable to press on the accelerator. In this context, there seem to be two readings of (27) (cf. Boylan 2023):

- (27) Brooke knows how to drive.
- a. \approx Brooke knows how *she* can drive. (*de se*; false)
 - b. \approx Brooke knows how *one* can drive. (generic; true)

One hypothesis, then, is that while the *de se* readings of deliberative knowledge require the subject to adopt a conditional intention, the generic readings do not. After all, it's hard to make sense of intending anything for a generic "one". In other words, while the *de se* reading of *S knows what to do* requires *S* to intend *for themselves* to do something, the generic reading doesn't require *S* to intend *for a generic person* to do something: it only requires *S* know *what one should do*.¹³ We might try to get at this difference by the following loose paraphrases:

- (28) *S* knows what to do.
- a. *de se*: *S* knows what *they* are to do.
 - b. generic: *S* knows what *one* is to do.

Thus, while Akira doesn't know what she is to do, since she doesn't intend *for herself* to study, Akira does know what one is to do, since she knows what *one* should do in her situation.

An alternative explanation, pertaining to the "can't bring herself to do it" reading, is that there might be ways to interpret the Akira case on which she is not akratic in the intended sense. The literature on akrasia generally recognizes two kinds. There is what we've been calling akrasia, which involves an agent deciding to do something they think they shouldn't do. But Holton (1999)

¹³ By "generic person", we do not necessarily mean an *average* person, but rather something more like an *ideally rational* person. Consider the following example (due to Una Stojnić (p.c.)): Imagine that Akira is a highly qualified surgeon and the only person in the world who can successfully carry out a certain surgery on her patient. Even if she suffers from akrasia, hesitating to perform the surgery because it's so dangerous, still *Akira knows what to do* seems true on the "generic" reading. Yet it's not true that a generic person should attempt the surgery since they would likely fail. Instead, what's true is that an ideal agent who found themselves in a similar situation to Akira with similar skillsets should attempt the surgery. This suggests the "generic" person is more of an idealization of Akira (e.g., stripping away her akratic tendencies) rather than an ordinary or average person.

distinguishes this kind of case from cases involving **weakness of will**, where an agent decides and intends to do something but fails to act on their intentions: these intentions fail to manifest in action somewhere downstream of the will. Akratic agents suffer from a conflict between their normative beliefs and their intentions. Weak-willed agents suffer from a conflict between their intentions and their actions.

When someone says, “Akira knows what to do. . . she just can’t bring herself to do it!” it’s plausible that they’re intending to convey Akira’s weakness of will, not akrasia proper. Akira knew she should study, and intended to study—but this intention was thwarted by the allure of the Xbox. This, we submit, is better described as a case of weakness of will.¹⁴

Granularity of Decisions. When describing the Parry case in section 3, we noted that there’s *some* sense in which Parry knows what to do: Parry knows to go to one of the parties. But in another sense, he fails to know what to do: he doesn’t know which party to go to.

We explain the difference between these two senses of *know what to do* using the decision-relativity of knowledge-to. There are two sorts of decisions Parry faces: a *coarse-grained* one and a *fine-grained* one. The former decision D_c only consists of two options: go to one of the parties and stay home. The latter decision D_f splits the first option into two parts: go to Ann’s party and go to Ben’s party. Our analysis predicts that Parry knows what to do in D_c (go to a party) but doesn’t know what to do in D_f (which party?).

What this shows is that knowledge-to ascriptions are highly sensitive to the way a decision is framed, in particular the granularity of the options involved. The more fine-grained the decision situation, the harder it is to have deliberative knowledge. Upon reflection, this makes sense. Most of the time, we do not make maximally specific plans for what to do: we first make a broad-level, coarse-grained decision and then refine it as needed.

Of course, it’s rare that we ever work out in full glory our plans of action. Even once Parry decides to go to Ann’s party, he may not have decided which route to take, or whether to grab his keys and then his wallet or the other way around. This doesn’t mean that we should always expect readings of knowledge-to ascriptions that are false. The level of granularity needed to evaluate a specific knowledge-to claim will be a highly context-sensitive matter, but it is unusual for speakers to have in mind decision situations that are needlessly fine-grained.

Moreover, some decision situations may contain certain “catch-all” options that obviate the need for finer granularity. For example, we might allow mixed strategies (e.g., Parry could decide by flipping a coin), or simply include an “arbitrarily pick” option. If “arbitrarily pick” is a genuine option for Parry and he decides to do this, then he could count as knowing what to do even if he has

¹⁴Of course, these distinctions are subtle. We suspect many English speakers will not be attuned to this when using or evaluating knowledge-to ascriptions. But once the distinction is drawn, we think weakness of will is a far more natural takeaway from (25).

yet to enact that option (cf. Ullmann-Margalit and Morgenbesser 1977). But if Parry is indecisive, going back and forth between going to Ann’s and going to Ben’s, then he still doesn’t count as knowing what to do, even if he knows he should just arbitrarily pick. In other words, arbitrarily deciding is compatible with deliberative knowledge. Indecision is not.

Should or Can? Recall from section 2 that knowledge-to ascriptions often seem paraphraseable using either *can* or *should* depending on certain factors, as illustrated by the following (repeated from (4)–(5)):

- (29) Magnus knows *where to get gas*.
 ≈ Magnus knows *where he/one can get gas*.
- (30) Hafdis knows *who to talk to*.
 ≈ Hafdis knows *who she/one should talk to*.

Of course, we think this “paraphrase” is acceptable only insofar as we assume the normative belief and intention align. But unless there’s good reason to think we’re dealing with a case of akrasia or parity, these sorts of paraphrases are generally harmless. One remaining question, however, is why each paraphrase is more acceptable than the other in certain cases.

Bhatt and Roberts both take a “*can*-first” approach by making the *can*-reading the primary one and deriving the *should*-reading under special conditions. For Bhatt, IQs contain a hidden modal $\diamond_{D,\rightarrow}$ that is interpreted roughly as “can and would satisfy the contextually relevant goals”. For Roberts, IQs concern what goals would be rational to pursue, where there’s no presumption that only one goal will be rational. On both these accounts, then, deliberative knowledge only strictly requires propositional knowledge of a *can*-claim. The *should*-readings arise when context makes it clear only one option can satisfy all the relevant goals or would be a rational goal to pursue.¹⁵

By contrast, our analysis of knowledge-to takes a “*should*-first” approach, making the *should*-reading the primary one. It’s just that the *should*-claim the subject must know could be disjunctive. Thus, Magnus knows where to get gas in virtue of knowing, say, that one should get gas from either Exxon or BP. The *can*-readings naturally arise from a free choice inference, like the following:

- (31) One should get gas at either Exxon or BP.
 ⇒ One can get gas at Exxon and one can get gas at BP.

Thus, by knowing one should get gas at either Exxon or BP, Magnus thereby knows one can get gas at Exxon and at BP.

On this view, the disjunctive *should*-claim is always available. However, in some conversational contexts, this *should*-claim is no more informative than the conjunctive *can*-claim, in which case the *can*-paraphrase will seem more natural. This is true especially in contexts where most of the available options are compatible with what the agent should do. For example, unless there are

¹⁵ See Hackl and Nissenbaum 2012 for criticism of Bhatt’s derivation of the *should*-reading.

huge differences in price, quality, or convenience, it doesn't really matter which station one gets gas from. Thus, the disjunctive *should*-claim *One should get gas at either Exxon or BP* is not much more informative than the conjunctive *can*-claim *One can get gas at Exxon and at BP*, which can explain why it's naturally paraphraseable as the latter.

One reason to prefer the *should*-first approach to the *can*-first approach is that in cases where multiple options are equally good, it seems incorrect to say one thereby knows to do each of those options, as illustrated by the following:

- (32) Magnus knows to get gas at either Exxon or BP.
 \Rightarrow Magnus knows to get gas at Exxon.
- (33) Parry knows to go to either Ann's party or Ben's party.
 \Rightarrow Parry knows to go to Ann's party.

This is hard to explain on the *can*-first approach. For one thing, it's well known that *can* admits of free choice inferences, as illustrated by (34)–(35).

- (34) One can get gas at either Exxon or BP.
 \Rightarrow One can get gas at Exxon and one can get gas at BP.
- (35) Parry is permitted to go to either Ann's party or Ben's party.
 \Rightarrow Parry is permitted to go to Ann's party and permitted to go to Ben's party.

So if the *can*-first approach were correct, i.e., if *S knows to ϕ* only requires *S* to know they can ϕ , we'd expect the inferences in (32)–(33) to sound plausible. Moreover, both Bhatt and Roberts directly predict that because Magnus knows one *can* get gas at Exxon, he knows *to* do this, since this is not a context where only one option would satisfy the relevant goals or be rational to pursue. Likewise, since Parry knows he *can* go to Ann's party, he know *to* do so. But this seems incorrect, especially if the options in question are exclusive.¹⁶

The *should*-first approach correctly predicts that knowing to ϕ or ψ does not imply knowing to ϕ . Magnus knows to get gas at Exxon only if he knows he should get gas there. But if getting gas at BP is equally permitted, it's not true that he *should* get gas at Exxon. Similarly, Parry knows to go to Ann's party only if he knows he should go to that party. But since he could just as well go to Ben's, it's not true he should go to Ann's party specifically.

Another benefit of the *should*-first approach is that it does a better job handling cases where an agent incorrectly believes certain options are compatible with what they should do. Suppose there are ten gas stations in town. As it so happens, the only station that's open is Exxon: the rest are closed. However, Magnus believes (with good evidence) that all ten stations are open, and so one can get gas at any of them. In that case, we hesitate to say that Magnus knows

¹⁶This problem is independent of the one we raised in section 3. Even if Bhatt and Roberts took a similar *should*-first approach, saying *S know what to do* is true iff *S* knows the relevant disjunctive *should*-claim, it would not be enough to address the cases of akrasia and parity. Both Akira and Parry already know the relevant disjunctive *should*-claim (though, in Akira's case, it's a "disjunction" with one disjunct), and yet neither knew what to do.

where to get gas simply because he has a correct belief that one can get gas at Exxon: anyone following his advice on where to get gas would fail 90% of the time! The *should*-first approach predicts this: to know where to get gas, Manus must know he/one should get gas from the open Exxon station, since that's the only option compatible with what he/one should do in this situation. The *can*-first approach, by contrast, predicts that Magnus does know where to get gas because he knows at least one location where he/one can get gas.

Still, the *can*-first approach does a better job at handling other cases, specifically those where the agent incorrectly believes certain options are *not* compatible with what they should do. Suppose Magnus knows he/one can get gas at Exxon, but falsely believes that the other nine stations are closed (in fact, they're all open and fully functional). In that case, it still seems right to say that Magnus knows where to get gas since he knows the Exxon station is open. It doesn't matter that he doesn't know the other stations are open: he just needs to know one place that's open.

We suspect this contrast can be explained by with the familiar distinction between two readings of a question: **mention-all** (exhaustive) readings and **mention-some** (non-exhaustive) readings. Mention-all questions require an exhaustive, complete answer whereas mention-some questions do not. One can answer the question *Where can one buy a pack of cigarettes?* just by naming one vendor that sells cigarettes, suggesting a mention-some reading. But to answer the question *Who is coming to the party?*, one needs to provide a complete list of attendees, suggesting a mention-all reading.

It's generally thought which reading is salient depends on context. For example, if a detective finds a special brand of cigarettes at the crime scene and wants to track down all the vendors that sell this brand, then a mention-all reading of *Where can one buy a pack of cigarettes?* becomes more salient. Or, if someone is just interested in learning whether they'll know someone at the party, a mention-some reading of *Who is coming to the party?* will be salient.

Throughout, we've been interpreting IQs as mention-all questions: to know what to do, the agent must know about *all* the options that are compatible with what they should do. This is why we require the agent to know a disjunctive *should*-claim. But IQs like *where to get gas* are naturally interpreted as mention-some questions: to know where to get gas, the agent only needs to know about *some* options compatible with what they should do. This is not mandatory, of course. Suppose Magnus is an event organizer and needs to know the location of every gas station in operation in order to direct people during the event. Magnus knows one can get gas at Exxon, but incorrectly believes the other stations are closed. In that context, it does not seem true that Magnus knows where to get gas, suggesting a mention-all reading (cf. Bhatt's observation that the context can affect whether an IQ has the force of a *should*-claim or a *could*-claim). Our hypothesis, then, is that whether a *should*-first or *can*-first approach is appropriate will depend on whether the embedded IQ is interpreted as a mention-all or mention-some question respectively.

A Propositional Reduction. We've argued that deliberative knowledge does not reduce to propositional knowledge. Our argument, in brief, is that deliberative knowledge requires the subject to form a conditional intention, whereas propositional knowledge does not.

One objection, however, is that *some* propositional knowledge does require subjects to form the relevant intention. An obvious example is knowledge *that* they form said intention. In other words, one might try to rescue the reduction of deliberative knowledge to propositional knowledge by strengthening our analysis so as to require not just that the subject *intend* to do such-and-such, but also to *know that they intend* such-and-such. For example, we might propose the following alternative analysis of *know to* constructions:

Propositional Analysis (*know to*)

S knows to ϕ in D iff (1) S knows they should ϕ in D , and (2) S knows they intend to ϕ in D .

This analysis yields the same results in the cases we've described so far: after all, if an agent is indecisive, then they can't know they've made a decision! But it succeeds at reducing knowledge-to to knowledge-that in a way that requires the agent to form a conditional intention.

How does this propositional analysis come apart from our non-propositional one? The answer depends on whether intention is a luminous mental state. If intention is luminous, then the two analyses will (almost) collapse.¹⁷ In that case, our analysis would predict that knowledge-to reduces to knowledge-that. If intention is not luminous, then we should expect that there will be cases where someone counts as knowing what to do on our analysis but not on this propositional analysis.

Unsurprisingly, we are not sympathetic with the idea that intention is a luminous mental state. Even if one is not convinced by Williamson's (2000) general argument against luminosity, we think it is independently implausible for intentions to be luminous: if only psychoanalysis were so simple! Still, we realize the issue is controversial and touches on the very nature of intention (cf. discussions of whether intentional action requires knowledge; e.g., Anscombe 1957; Davidson 1973, 1978; Moran 2001; Newstead 2006; Paul 2009, 2011; Schwenkler 2015; Piñeros Glasscock 2020; Beddor and Pavese 2022; Shepherd and Carter 2022). If intention is luminous, then we must concede that knowledge-to does indeed reduce to knowledge-that (or at least to being in a position to possess knowledge-that). Even so, we maintain that our account is an improvement over extant accounts in the literature, which do not recognize a requirement for the subject of knowledge-to ascription to form a conditional intention, in the event that this comes apart from possessing the corresponding normative knowledge.

If intention is not luminous, and so the two analyses come apart, we suspect ours fares better. Imagine that Akira falsely thinks she is akratic. So she knows

¹⁷ We say "almost" since a luminous condition only requires that one is *in a position to know* that it obtains, not necessarily that one *knows* it obtains, if one is in it.

that she should study and she *incorrectly thinks* she's indecisive, when really, she intends to study. In that case, we think the following sounds plausible:

(36) Akira knows what to do. She just doesn't realize it (yet).

The propositional analysis, by contrast, predicts that (36) is false since Akira does not know (or even believe) she intends to study. We think this prediction is incorrect: Akira can know what to do "implicitly", as it were.

In short, then, we think our analysis of deliberative knowledge, to use game-theoretic terminology, "strictly dominates" the propositional analysis: the propositional analysis fares no better than ours if intention is luminous, and it fares worse if not.

6 Open Questions

In closing, we discuss some open questions about deliberative knowledge.

Other IQ-Embedding Constructions. Throughout, we've focused exclusively on *know*+IQ constructions. But other verbs that take interrogative complements also take IQ-complements. Bhatt (1999, pp. 115–118) lists several classes of IQ-accepting verbs, including:

- verbs of knowledge: *forget, remember, learn, figure out*
- verbs of decision: *decide, choose, determine, control*
- verbs of communication: *explain, debate, show, tell*
- verbs of inquiry: *consider, ask, wonder*
- select verbs with *about* or *over*: *think about, divided over, unsure about*.

Bhatt also observes that some verbs that take interrogative complements do not take IQ-complements, such as:

- verbs of conjecture: *predict, guess, estimate*
- verbs of relevance: *matter, be important, care*
- verbs of dependency: *depend on, have an influence on, make a difference to*
- emotive predicates: *surprising, amazing, appalling*

Given this, it's natural to wonder whether our analysis of *know*+IQ can be generalized to other IQ-embedding constructions. To do this, we would ideally like a fully compositional semantics for IQs so that we can decompose the meaning of verb+IQ into the semantics of the verb and the semantics of IQs. We suspect this can be done, but we do not have strong views about how exactly to develop a fully compositional story for all these constructions.¹⁸

¹⁸ A natural suggestion we will simply float is to generalize Santorio's (2016) expressivist semantics for *know how* to *know*+IQ more generally. So we assign IQs to partitions on sets of world-plan pairs. (Santorio uses "performance plans", which need not involve forming intentions or being manifested in dispositions to act. The plans we would need to analyze other IQs would have to be something closer to plans of the ordinary variety.) While Santorio bases his approach on Bhatt's, we could alternatively combine Santorio's expressivist semantics with Roberts's analysis of IQs by assigning them to functions from agents to partitions on world-plan pairs.

One complication is that the list of verbs that can take IQ-complements does not neatly align with the list of verbs that can take infinitival complements. Not all IQ-accepting verbs accept infinitival clauses. For example, *forget to ϕ* , *remember to ϕ* , *learn to ϕ* are all grammatical, while *control to ϕ* , *explain to ϕ* , *consider to ϕ* , *wonder to ϕ* , and *think about to ϕ* are not. These don't line up with the verbs that can take that-clauses: e.g., *explain that ϕ* and *consider that ϕ* are fine (this could be seen as another argument against propositional approaches). Conversely, some of the verbs that can't take IQs *can* take infinitival clauses. Thus, while *predict to ϕ* , *estimate to ϕ* , *matter to ϕ* , and *depend on to ϕ* are marked,¹⁹ *be important to ϕ* and *care to ϕ* are acceptable. This means we may not always be able to bootstrap an analysis of verb+IQ from verb+to ϕ .

Furthermore, these verbs do not all require the same sorts of conditions we've argued that *know*+IQ requires. For example, while one must form an intention to *know* what to do, one does not need to form an intention to *consider*, *ask*, or *debate* what to do. Similarly, if one *forgets* what to do, one doesn't thereby intend to do that which was forgotten—quite the opposite! On the other hand, while one must know what one should do to know what to do, one does not need to *decide* what one should do in order to decide what to do (esp. if akrasia is possible). Indeed, the literature on deliberative questions suggests one can also consider, ask, wonder, and debate what to do without considering, asking, wondering, and debating what one should do, and vice versa.

Finally, it's natural to think that the content of an IQ can just be analyzed as a set of options, or actions, that a subject can take, or perhaps more generally as a function from subjects, decision situations, and worlds to sets of options. This is tricky, however, in light of certain coordination data suggesting arguments of similar types:

- (37) Tanya knows that there's a housing crisis and who to blame for it.
 (38) Connor considered who he can call and whether to call them.

As Roberts points out, however, (following Dowty and Jacobson 1991), the data does not obviously require analyzing the coordinated constituents as having the same semantic type. One reason for this is that we see the same coordination data between (even finite) interrogatives and noun phrases (Sag, Gazdar, Wasow, and Weisler, 1985):

- (39) Rob knows where the grocery store is and the fastest route there.
 (40) Priscilla predicted when the stock market would crash and the reason it did.

Thus (to summarize with an example of the very phenomenon), it is still contentious what to make of this data and what the best account of it is.

¹⁹ These are grammatical when the infinitive results from movement outside the scope of verb (see footnote 8). For example, *What do you predict to see?* or *The number of stocks I estimate to crater in the next several days is high* are fine.

Knowledge-how. Earlier, when we defined deliberative knowledge, we set aside IQs involving *how to* questions. Our reason for this was dialectical: we wanted to show that independently of the relation between know-how and abilities, knowledge-to is not reducible to knowledge-that. Still, it is natural to ask whether our analysis of knowledge-to bears on the debates over knowledge-how.

On the one hand, it would be surprising if *how to* questions behaved radically differently from other IQs. Ideally, one might think, it would be desirable to give a unified treatment of all IQs. We could, for instance, postulate that *how to* questions restrict the options in a decision situation to those of the form:

how to ϕ : ϕ by w_ϕ , where w_ϕ is a way to ϕ

On the other hand, there are reasons to think *how to* questions are distinctive enough to deserve a separate treatment. Specifically, there seems to be a difference between practical and deliberative knowledge in terms of what the agent intends.²⁰ To see this, observe first that the following inference fails:

- (41) Sandy knows how to make a sandwich.
 \Rightarrow Sandy intends to make a sandwich.

Now, this alone does not distinguish practical and deliberative knowledge. As we saw earlier, the corresponding inference for other IQs also fails, as our account correctly predicts:

- (42) Sandy knows where to get a sandwich.
 \Rightarrow Sandy intends to get a sandwich.

The IQ *where to get a sandwich* restricts the decision situation D to (positive boolean combinations of) options of the form *get a sandwich at l* . Similarly, one might say that *how to make a sandwich* restricts the decision situation to (positive boolean combinations of) options of the form *make a sandwich by w_s* . So while knowing how to ϕ doesn't require an intention to ϕ full stop, it would require an intention to ϕ relative to a restricted decision where the options all involve ϕ -ing in some way or other.

²⁰ Another difference is that knowledge-how ascriptions can be put into comparative forms, as illustrated below (Sgaravatti and Zardini, 2008; Bengson and Moffett, 2011; Michaelis, 2011; Wiggins, 2012; Pavese, 2017):

- (i) a. Skylar knows how to ski better than they know how to snowboard.
 b. Skylar knows how to ski better than Steve.

However, such comparative constructions using deliberative knowledge ascriptions are not acceptable (cf. Roberts's (2009) example involving *whether*):

- (ii) a. #Akira knows to study better than to play video games/better than Beth.
 b. #Parry knows what to wear better than where to go/better than Kenneth.
 c. #Magnus knows where to get gas better than when/better than Miranda.

This suggests that comparative knowledge-how ascriptions require a different kind of analysis.

Even so, there are reasons to think practical knowledge does not require the subject to adopt *any* intentions, even with respect to these restricted decisions. Suppose Parry decides to go to Ann's party. The question Parry now faces is how to get there. He knows two different routes but hasn't decided which to take. Here, there seems to be no sense in which Parry does not know how to get to Ann's party. To be sure, he hasn't *decided* how to get to Ann's party, but he *knows* how to do so. In other words, the following inference seems patently bad:

- (43) Parry hasn't decided how to get to Ann's party.
⇒ Parry doesn't know how to get to Ann's party.

By contrast, the following seem fine:

- (44) Parry hasn't decided which party to go to.
⇒ Parry doesn't know which party to go to.
- (45) Parry hasn't decided when to go to Ann's party.
⇒ Parry doesn't know when to go to Ann's party.

So while indecision precludes deliberative knowledge, it does not seem to preclude practical knowledge.

Furthermore, knowledge-how arguably plays a special role in deliberation: practical knowledge seems to be a precondition for an action to count as a genuine option for an agent in the first place. We might articulate this in terms of the following principle (cf. Hornsby 2016; Habgood-Coote 2017; Beddor and Pavese 2022):

Options Require Knowledge-how

If D is a (possibly hypothetical) decision situation for S and ϕ is an option in D , then S is faced with D only if S knows how to ϕ .

Note, this principle does not say that one cannot consider, deliberate about, or form intentions about what to do in a decision situation D without knowing how to do each option in D . One may think about a purely hypothetical decision concerning whether to go to Mars even if one currently doesn't have any idea how to get there. Rather, it says that in order for an agent to *find themselves* faced with a decision, each option must be something the agent knows how to do.

This principle is not unassailable. You could plausibly decide to write a book even without knowing how to do that. A proponent of this principle might reply that in such a case, you do not have the option to write a book, but rather to take steps towards writing a book (Habgood-Coote, 2017), or to try to write a book (Holguín and Lederman, 2023). The plausibility of this principle is thus very sensitive to how the options in a decision situation are conceived.

But whatever we make of this principle, there seems to be no plausible analogue governing other *know*+IQ constructions. One does not need to know when to ϕ , where to ϕ , and so on in order for ϕ -ing is a genuine option. Facing a decision, in other words, does not require having already decided on the exact

time, location, reason, and so forth to perform each action. Figuring all that out is precisely what deliberating is about!

On both philosophical and linguistic grounds, then, there are reasons to think practical knowledge deserves a unique analysis independent of the analysis of deliberative knowledge. Nothing we've said in favor of a non-propositional analysis of other *know*+IQ rules out a strictly propositional, intellectualist analysis of *know how* on which it reduces to *know that*.

Knowledge and Virtue. Socrates famously claimed that virtue is a kind of knowledge (Protagoras 352c). The oldest argument against this "Socratic intellectualism" about virtue is the possibility of akrasia. Characters like Akira, who know they should do one thing but choose to do another, seem to be direct counterexamples.²¹ This is partly why akrasia has been so hotly contested.

We think our theory of deliberative knowledge can be used to chart a middle course in this debate, one that preserves the spirit of Socratic intellectualism while acknowledging the possibility of akrasia. An implicit assumption in the argument from akrasia is that the kind of knowledge that virtue could be for Socrates is propositional. Thus the need for Socrates to deny the possibility of akrasia: such agents would have propositional knowledge of the good but lack virtue. If we are right, though, not all knowledge is propositional. Deliberative knowledge *constitutively* involves the formation of apt intentions—intentions, that is, to do things consistent with what one should do.

We are inclined, therefore, to accept:

Deliberative Socratic Intellectualism (DSI)

A virtuous agent (in some domain) is one who knows what to do (in that domain).

DSI, we think, can maintain the benefits of traditional Socratic intellectualism while avoiding its pitfalls (cf. Segvic's (2000) interpretation of Socrates). In particular, worries about the possibility of akrasia do not get off the ground, for agents with deliberative knowledge are precisely those agents who have formed their intentions in accordance with their knowledge of normative truth. Agents with deliberative knowledge are neither akratic nor paralyzed by parity.

The most attractive feature of Socratic intellectualism, at least for Socrates, is that it implies virtue can be taught. Can, then, deliberative knowledge be taught? We think it can. At least in English, 'taught' can take both IQs and infinitives:

- (46) The Native Americans taught the settlers to harvest corn before the winter.
- (47) Akira's tutor taught her what to study for the exam and how much time to allocate to studying.

²¹ There's some historical debate over the extent to which Socrates held akrasia to be impossible. For discussion, see Devereux 1995; Segvic 2000; Brickhouse and Smith 2010; Weiss 2006; Kamtekar 2017.

- (48) Calculus teachers teach students to use L'Hôpital's rule for limits with indeterminate forms.

Independently of English data, we find it plausible that intentions can be acquired through teaching. The pupils in these examples are ones who, as a result of a process of teaching, come to have deliberative knowledge about some domain both by being taught what is good and being taught to act in accordance with it. Thus, if DSI is right, these are agents who have been taught to be virtuous. The core motivation for Socratic intellectualism is therefore maintained, while jettisoning the problematic assumption (propositionalism) that led to refutations by akrasia and parity.

Cross-linguistic Data and the Philosophical Significance of the Deliberative-Normative Divide. One final open question about deliberative knowledge ascriptions is their cross-linguistic availability. It has been noted in the literature on practical knowledge that not all languages express knowledge-how ascriptions in the exact same way (Rumfitt, 2003; Wiggins, 2012; Abbott, 2013; Ditter, 2016; Hornsby, 2016; Pavese, 2021). The same is true of *know*+IQ constructions generally. For one, not every language permits one to embed an infinitival clause under the corresponding word for *know*. In Turkish, for example, one must use the finite verb form.²² In Hungarian, one must use either a deontic modal or an imperative.²³ Some languages do not even have an infinitival verb form, such as modern Greek (to ascribe practical knowledge, one must instead use the first-person verb form). Even among those that do permit one to embed an infinitival clause, it is not interpreted the same way as in English. In Spanish, French, Italian, German, and Russian, for example, the *know to* construction is interpreted as knowledge-how or ability. In both German and Russian, there is arguably even a different word for *know how* than for *know that*.

What does this cross-linguistic variability mean for deliberative knowledge? One thing it might suggest is that deliberative knowledge is not a distinctive *kind* of knowledge. If not all languages permit the *know*+IQ construction, or interpret such constructions the same way, one might reasonably conclude that “deliberative knowledge” as we call it is really just a shorthand for ordinary propositional knowledge plus intention. In a way, our analysis of these constructions in English agrees with this: after all, we argue that we *can* reduce knowledge-to to a combination of (normative) knowledge-that and intention.²⁴ Of course, this is not a “propositional reduction”, in that knowledge-to is not completely reducible to knowledge-that. But it is reductive in that it does not simply leave knowledge-to unanalyzed.

²² Thanks to Zeynep Soysal for clarifying this to us.

²³ Thanks to Zoltán Szabó for pointing this out to us.

²⁴ Interestingly, the cross-linguistic argument is often used in the opposite way in the case of knowledge-how: it's argued that *because* different languages express knowledge-how differently (e.g., not always using the same verb for *know that* and *know how*), that suggests knowledge-how is not reducible to knowledge-that.

Does this mean that the distinction between deliberative and normative questions is not philosophically substantive? Are the supporters of the deliberative-normative divide simply caught up in a quirk of English? Not necessarily. Just because some languages do not have *know*+IQ constructions, or they have such constructions but do not interpret them in the same way as English, that doesn't mean other languages have no way of expressing deliberative knowledge ascriptions. Indeed, if our analysis is correct, other languages can translate such ascriptions as conjunctions of knowledge-that and intention ascriptions. Similarly, the fact that other languages do not represent the difference between deliberative and normative questions grammatically in the same way as English does not imply that speakers of other languages cannot recognize the difference. The distinction between deliberative and normative questions is ultimately a conceptual distinction. Nothing requires this distinction be reflected in natural language constructions.

Still, this cross-linguistic variability does suggest that English-speaking philosophers should be cautious about using linguistic arguments to draw philosophical conclusions about the relation between deliberative and normative questions. There might be a neat way to express this distinction in English, but the fact we can state a distinction doesn't make it a real or helpful distinction. The same lesson has been drawn by intellectualists in the literature on practical knowledge: just because other languages encode the distinction between propositional and practical knowledge linguistically, it doesn't follow that they are separate kinds of knowledge. The legitimacy of the distinction between deliberative and normative questions ought to be grounded in more general philosophical and empirical considerations concerning the relation between belief, intention, and normativity.

References

- Abbott, Barbara. 2013. "Linguistic Solutions to Philosophical Problems: The Case of Knowing How." *Philosophical Perspectives* 27:1–21.
- Anscombe, Elizabeth. 1957. *Intention*. Oxford: Basil Blackwell Publishers.
- Balcerak Jackson, Brendan. 2019. "Essentially Practical Questions." *Analytic Philosophy* 60:1–26.
- Beddor, Bob and Pavese, Carlotta. 2022. "Practical Knowledge without Luminosity." *Mind* 131:917–934.
- Bengson, John and Moffett, Marc. 2011. "Non-propositional Intellectualism." In John Bengson and Marc Moffett (eds.), *Knowing How: Essays on Knowledge, Mind, and Action*, 161–195. Oxford: Oxford University Press.
- Bhatt, Rajesh. 1999. *Covert Modality in Non-Finite Contexts*. Ph.D. thesis, University of Pennsylvania. Published in 2006. Berlin: de Gruyter Mouton.

- Boylan, David. 2023. "Knowing How is Knowing How You Are (or Could Have Been) Able." Manuscript.
- Bratman, Michael. 1987. *Intention, Plans, and Practical Reason*. Cambridge, MA: Harvard University Press.
- Braun, David. 2006. "Now You Know Who Hong Oak Yun Is." *Philosophical Issues* 16:24–42.
- Brickhouse, Thomas C and Smith, Nicholas D. 2010. *Socratic Moral Psychology*. Cambridge: Cambridge University Press.
- Broome, John. 2008. "Comments on Allan Gibbard's Tanner Lectures." In Barry Stroud (ed.), *Reconciling Our Aims: In Search of Bases for Ethics*, 102–119. Oxford: Oxford University Press.
- Cariani, Fabrizio, Kaufmann, Magdalena, and Kaufmann, Stefan. 2013. "Deliberative modality under epistemic uncertainty." *Linguistics and Philosophy* 36:225–259.
- Chang, Ruth. 2002. "The Possibility of Parity." *Ethics* 112:659–688.
- Chislenko, Eugene. 2016. "A Solution for Buridan's Ass." *Ethics* 126:286–310.
- Ciardelli, Ivano, Groenendijk, Jeroen, and Roelofsen, Floris (eds.). 2018. *Inquisitive Semantics*. Oxford: Oxford University Press.
- Clarke-Doane, Justin. 2020. *Morality and Mathematics*. Oxford: Oxford University Press.
- . 2022. "From Non-Usability to Non-Factualism." *Analysis* 81:747–758.
- Constantin, Jan. 2018. "A dispositional account of practical knowledge." *Philosophical Studies* 175:2309–2329.
- Davidson, Donald. 1973. "Freedom to act." In Ted Honderich (ed.), *Essays on Freedom and Action*. London: Routledge and Kegan Paul.
- . 1978. "Intending." *Philosophy of History and Action* 11:41–60.
- . 1980. "How is Weakness of the Will Possible?" In *Essays on Actions and Events*, 21–42. Oxford: Oxford University Press.
- Devereux, Daniel. 1995. "Socrates' Kantian Conception of Virtue." *Journal of the History of Philosophy* 33:381–408.
- Ditter, Andreas. 2016. "Why Intellectualism Still Fails." *The Philosophical Quarterly* 66:500–515.
- Dowty, David and Jacobson, Pauline. 1991. "Infinitival questions." Unpublished manuscript.

- Faraci, David. 2017. "On Leaving Room for Doubt: Using Frege-Geach to Illuminate Expressivism's Problem with Objectivity." In Russ Shafer-Landau (ed.), *Oxford Studies in Metaethics*, 244–264. Oxford: Oxford University Press.
- Farkas, Katalin. 2017. "Practical Know-Wh." *Noûs* 51:855–870.
- George, B. R. 2013. "Knowing-'wh', Mention-Some Readings, and Non-Reducibility." *Thought* 2:166–177.
- Gibbard, Allan. 2003. *Thinking How to Live*. Cambridge, MA: Harvard University Press.
- . 2008. *Reconciling Our Aims: In Search of Bases for Ethics*. Oxford: Oxford University Press.
- Groenendijk, Jeroen and Stokhof, Martin. 1982. "Semantic Analysis of Wh-Complements." *Linguistics and Philosophy* 5:172–233.
- Habgood-Coote, Joshua. 2017. *Knowledge-How: Linguistic and Philosophical Considerations*. Ph.D. thesis, University of St. Andrews.
- Hackl, Martin and Nissenbaum, Jon. 2012. "A modal ambiguity in *for*-infinitival relative clauses." *Natural Language Semantics* 20:59–81.
- Hamblin, Charles Leonard. 1958. "Questions." *Australasian Journal of Philosophy* 36:159–168.
- . 1973. "Questions in Montague English." *Foundations of Language* 10:41–53.
- Hare, Richard M. 1952. *The Language of Morals*. Oxford: Oxford University Press.
- Hedden, Brian. 2012. "Options and the subjective *ought*." *Philosophical Studies* 158:343–360.
- Hieronymi, Pamela. 2009. "The Will as Reason." *Philosophical Perspectives* 23:201–220.
- . 2011. "Reasons for Action." *Proceedings of the Aristotelian Society* 111:407–427.
- Holguín, Ben and Lederman, Harvey. 2023. "Trying without fail." Manuscript.
- Holton, Richard. 1999. "Intention and Weakness of Will." *Journal of Philosophy* 96:241–262.
- Hornsby, Jennifer. 2016. "Intending, knowing how, infinitives." *Canadian Journal of Philosophy* 46:1–17.
- Kamtekar, Rachana. 2017. *Plato's Moral Psychology: Intellectualism, the Divided Soul, and the Desire for Good*. Oxford: Oxford University Press.

- Karttunen, Lauri. 1977. "Syntax and Semantics of Questions." *Linguistics and Philosophy* 1:3–44.
- Lewis, David K. 1988a. "Relevant Implication." *Theoria* 54:161–174.
- . 1988b. "Statements Partly About Observation." *Philosophical Papers* 17:1–31.
- Masto, Meghan. 2010. "Questions, answers, and knowledge-wh." *Philosophical Studies* 147:395–413.
- Michaelis, Laura. 2011. "Knowledge Ascription by Grammatical Construction." In John Bengson and Marc Moffett (eds.), *Knowing How: Essays on Knowledge, Mind, and Action*, 261–282. Oxford: Oxford University Press.
- Moran, Richard. 2001. *Authority and Estrangement*. Princeton: Princeton University Press.
- Newstead, Anne. 2006. "Knowledge by Intention? On the Possibility of Agent's Knowledge." In Stephen Hetherington (ed.), *Aspects of Knowing*, 183–201. Amsterdam: Elsevier Science Publishers.
- Noë, Alva. 2005. "Against Intellectualism." *Analysis* 65:278–290.
- Owens, David. 2011. "Deliberation and the First Person." In Anthony Hatzimoysis (ed.), *Self Knowledge*, 261–277. Oxford: Oxford University Press.
- Paul, Sarah K. 2009. "How We Know What We're Doing." *Philosophers' Imprint* 9:1–24.
- . 2011. "How we know what we intend." *Philosophical Studies* 161:327–346.
- Pavese, Carlotta. 2015. "Practical Senses." *Philosophers' Imprint* 15:1–25.
- . 2017. "Know-How and Gradability." *The Philosophical Review* 126:345–383.
- . 2021. "Knowledge How." In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*. Palo Alto, CA: Metaphysics Research Lab, Stanford University.
- Piñeros Glasscock, Juan S. 2020. "Practical Knowledge and Luminosity." *Mind* 129:1237–1267.
- Risberg, Olle. 2023. "Ethics and the Question of What to Do." *Journal of Ethics and Social Philosophy* 25:376–412.
- Roberts, Craige. 2009. "know-how: A Compositional Approach." In Erhard Hinrichs and John Nerbonne (eds.), *Theory and Evidence in Semantics*, 183–213. CSLI Publications.
- Rumfitt, Ian. 2003. "Savoir Faire." *Journal of Philosophy* 100:158–166.
- Ryle, Gilbert. 1949. *The Concept of Mind*. London: Hutchinson.

- Sag, Ivan A, Gazdar, Gerald, Wasow, Thomas, and Weisler, Steven. 1985. "Coordination and how to distinguish categories." *Natural Language & Linguistic Theory* 3:117–172.
- Santorio, Paolo. 2016. "Nonfactual Know-How and the Boundaries of Semantics." *The Philosophical Review* 125:35–82.
- Schaffer, Jonathan. 2007. "Knowing the Answer." *Philosophy and Phenomenological Research* 75:383–403.
- Schwarz, Wolfgang. 2021. "Objects of Choice." *Mind* 130:165–197.
- Schwenkler, John. 2015. "Understanding "Practical Knowledge"." *Philosophers' Imprint* 15:1–32.
- Segvic, Heda. 2000. "No One Errs Willingly: The Meaning of Socratic Intellectualism." *Oxford Studies in Ancient Philosophy* 19:1–45.
- Setiya, Kieran. 2008. "Practical Knowledge." *Ethics* 118:388–409.
- . 2012. "Knowing How." *Proceedings of the Aristotelian Society* 112:285–307.
- Sgaravatti, Daniele and Zardini, Elia. 2008. "Knowing How to Establish Intellectualism." *Grazer Philosophische Studien* 77:217–261.
- Shepherd, Joshua and Carter, J Adam. 2022. "Knowledge, Practical Knowledge, and Intentional Action." *Ergo* 9:556–583.
- Southwood, Nicholas. 2016. "'The Thing to Do' Implies 'Can'." *Noûs* 50:61–72.
- Stanley, Jason and Williamson, Timothy. 2001. "Knowing How." *Journal of Philosophy* 98:411–444.
- Ullmann-Margalit, Edna and Morgenbesser, Sidney. 1977. "Picking and Choosing." *Social Research* 44:757–785.
- Weirich, Paul. 1983. "A Decision Maker's Options." *Philosophical Studies* 44:175–186.
- Weiss, Roslyn. 2006. *The Socratic Paradox and Its Enemies*. Chicago: University of Chicago Press.
- Wiggins, David. 2012. "Practical Knowledge: Knowing How To and Knowing That." *Mind* 121:97–130.
- Williamson, Timothy. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.