



Knowing more (about questions)

Joshua Habgood-Coote¹ 

Received: 2 June 2021 / Accepted: 25 January 2022

© The Author(s) 2022

Abstract

How should we measure knowledge? According to the *Counting Approach*, we can measure knowledge by counting pieces of knowledge. Versions of the Counting Approach that try to measure knowledge by counting true beliefs with suitable support or by counting propositions known run into problems, stemming from infinite numbers of propositions and beliefs, difficulties in individuating propositions and beliefs, and cases in which knowing the same number of propositions contributes differently to knowledge. In this paper I develop a novel question-relative and contextualist version of the counting approach, which measures an agent's knowledge by counting the number of complete answers of a contextually salient issue they can rule out. The question-relative and contextualist version of the Counting Approach avoids the issues for the proposition and belief-based systems, and offers a general, systematic, and explanatory system for measuring knowledge.

Keywords Knowledge · Questions · Subject-matter · Measurement theory · Inquiry

1 Introduction

What is it for one person to know more than another? At first pass, we might have thought that knowing more is a matter of knowing more things; a matter of having more pieces of knowledge. Let's call the view that we can make sense of the measure of knowledge in terms of countable pieces of knowledge the counting approach. We can individuate pieces of knowledge by referring either to the relation or to the relation of knowledge, suggesting two versions of the counting approach:

COUNTING BELIEFS: A knows more than B just in case the number of A's justified true beliefs is larger than the number of B's justified true beliefs.

✉ Joshua Habgood-Coote
j.habgood-coote@leeds.ac.uk

¹ School of Philosophy, Religion, and History of Science, University of Leeds, Leeds, UK

COUNTING PROPOSITIONS: A knows more than B just in case the number of propositions which A knows is larger than the number of propositions that B knows.

An account of what it is to know more will need to do substantive philosophical work elsewhere in philosophy. An account of the measurement of knowledge will help us to weigh up different projects of inquiry (Treanor, 2014), to understand expertise in terms of knowing more (Goldman, 2001), and to make sense of the connection between scientific progress and increases in knowledge (Bird, 2007).

In a series of papers, Treanor (2013, 2014, 2019) has presented a number of objections to both versions of the counting approach.¹ Against COUNTING BELIEFS, Treanor points out that it is a live option that all agents have an infinite number of beliefs meeting the relevant standard (2013, p. 580), observes that we do not have a clear principle of individuation for beliefs (2013, pp. 582–587), and gives several cases in which gaining the same number of knowledge-worthy beliefs seems to contribute differentially to knowledge (2013, pp. 588–591).² Similar points apply to COUNTING PROPOSITIONS: it is a live option that we know an infinite number of propositions (or at least that we are ignorant of infinitely many propositions) (2013, p. 581, n.d.), the principle of individuation for propositions is no clearer than that for beliefs (2013, pp. 586–587, n.d.),³ and the same cases demonstrate that learning the same number of propositions can contribute differentially to knowledge.⁴

My goal in this paper is develop a novel version of the counting approach which measures knowledge by counting *answers to questions*.⁵ This approach endorses the following account of knowing more about a question:

COUNTING ANSWERS: A knows more than B about a question Q just in case A's knowledge rules out more of the possible complete answers to Q than B's knowledge does.

This approach has three components.

¹ Treanor (2013) doesn't sharply distinguish between COUNTING BELIEFS and COUNTING PROPOSITIONS, but Treanor (2019) focuses on COUNTING PROPOSITIONS.

² These cases are important for understanding trivial truths cases (Treanor, 2013, pp. 598–9, 2014).

³ We can individuate propositions by considering the worlds they are true in, but this will get into trouble with necessarily true propositions. Another option would be to appeal to atomic propositions, but this would take a considerable amount of working out.

⁴ Treanor canvasses two alternatives to the counting approach: a modal approach that measures an agent's knowledge in terms of the counterfactual distance between the actual world and the furthest world which is compatible with her knowledge, and a similarity-based approach that reduces how much agents know to how similar their representational states are to the world. Treanor cautiously endorses the similarity-based approach, due to the obscurity of the notion of counterfactual distance.

⁵ 'Question' is multiply ambiguous, referring to entities, a kind of sentence or clause, and a kind of speech act. As I will use the term, a question is an entity on par with propositions that provides the semantic value for interrogative clauses, and can be expressed via the speech act of asking. I italicise questions (and propositions), and put interrogatives (and other sentences) in single quotes.

The first is an argument that the interesting phenomenon in comparative sentences is not knowledge itself, but knowledge about issues.^{6,7} I will use ‘issue’ as a term of art to refer to the genus of which questions and subject-matters are species, and the final statement of COUNTING ANSWERS will invoke issues. Although not often discussed (but see Yablo, 2014) issues are the kinds of things which mental states, conversations, and books are about, which experts are experts in, and what understanding is of. Whereas a proposition corresponds to a particular way the world might be, an issue corresponds to a set of ways the world—or some part of it—might be. I will spend more time developing this idea in the case of questions, because it is intuitively easier to grasp some of the core ideas, but the full account claims that we should measure knowledge in an issue-relative way.

The second component is a context-sensitive semantics for sentences comparing knowledge, which allows that the questions against which knowledge is measured can vary depending on context. This semantics relies on context-sensitivity relating to the questions and subject-matters which provide the background for measurement. This idea is orthogonal to whether ‘knows’ expresses different relations in different contexts (Cohen, 1986; DeRose, 2009; Lewis, 1996): the idea is not that the standards for knowledge shift, but that the scale for measuring amounts of knowledge is provided by context. For simplicity I will assume a noncontextualist view of ‘knows’. Below we will discuss contextualist, invariantist, and subject-sensitive views of the origins of questions in comparative sentences below, but these views are closer to contextualism and invariantism about the questions in knowledge-wh ascriptions (see Braun, 2006, 2011; Parent, 2014) than the similarly named views about knowledge-that ascriptions.

The final component is an account of the abstract structure of issues more generally. I will employ a partition-based account of questions and subject-matters, which allows us to measure knowledge by counting how many of the cells in the relevant partition are ruled out by agents’ knowledge.⁸

⁶ For official purposes, I will assume that knowledge (picked out by both knows-that and knows-wh ascriptions) is a binary relation between agents and standard propositions. There are various ways to involve issues in knowledge (either -that or -wh), any of which would also be compatible with COUNTING ANSWERS: (i) knowledge might be a three-place relation between agents, propositions, and questions (Schaffer, 2005, 2008) (see also Yalcin (2018)), (ii) knowledge might be a relation to directed propositions that combine a traditional proposition and question (Pavese, n.d.; Yablo, n.d., 2014, Chapter 7), (iii) knowledge might be a two-place relation between agents and questions (Habgood-Coote, 2018; Mastro, 2010), or iv) ‘knows’ might be ambiguous between question and proposition-relating senses (Karttunen 1977). Although COUNTING ANSWERS is a cousin of contrastivist views of knowledge, it does not entail this view: knowledge might be a binary relation to propositions which is best measured by questions.

⁷ One might want to distinguish between knowing more about an issue, knowing more of an issue, and knowing more relative to an issue. Ordinary language is slippery here, but it is important to distinguish the ability to rule out more answers (which I will call knowing more about an issue), knowing what the possibilities in an issue are without being able to rule them out (which we might call knowing more of an issue), and knowing more about both an issue and relevant other issues (which we might call knowing more relative to an issue). Our interest is in the first sense, and I will use ‘knowing more about’ to pick it out.

⁸ Other approaches to questions are available, identifying them with unique true answers, sets of possible answers, or sets of resolution conditions (for an overview, see Roelofsens, 2019). I will work within the partition-based account as it gives us all the resources we need, whilst remaining fairly straightforward. We could get similar results (with some extra leg work) if we switched to the inquisitive semantics approach of identifying questions with resolution conditions (Ciardelli et al., 2018), or identified questions with sets of direct answers (Wisniewski 1995).

The plan of action is as follows. Section 2 makes some comments about measurement systems and the difficulties with counting knowledge. Sections 3 and 4 build an argument that the interesting comparisons of knowledge states are question-relative, appealing to ordinary language, and the relation between knowledge and inquiry. Section 5 gives the details of the account of questions, and section 6 presents the account. Section 7 deals with some loose ends.

2 Measurement

A system of measurement aspires to be a mapping from a worldly phenomenon to some abstract structure that preserves some of the important features of the worldly phenomenon. The ideal for a measurement system is a mathematical measurement theory which captures a homomorphism between the objects and the abstract objects (success at this task being proved by establishing a representation theorem for that system). For example, the metric system of measurement for weights provides a mapping from the masses of physical objects to natural numbers, preserving the order of the masses of the objects in the physical system as well as relations of equality, and proportions between masses.⁹

There are lots of ways to measure stuff. When dealing with numerical measurement systems, we can obtain a different system by multiplying the numbers by a constant factor, or by shifting the zero point on the scale (think about the difference between Celsius and Fahrenheit). More interestingly, we might change the abstracta being appealed to (using sets rather than numbers), or appeal to different mappings.

There may be several ways to measure knowledge that are equally legitimate, which might turn out to be useful for different purposes. The goal of this paper is not to argue that COUNTING ANSWERS is the only correct measurement system for knowledge—I don't think that there is any such thing—but to argue that it is a good measurement system with some attractive features.

What makes a good measurement system? I suggest that we want a system that is general, systematic, and explanatory.

Generality requires that a measurement system maps all the target phenomena to an element in the abstract system. A system for measuring masses would be inadequate if it assigned numerical values only to objects with masses of less than a tonne.

Systematicity requires that the mapping from the physical objects to abstracta has some principled basis, being generated from some uniform function rather than being ascribed on an ad hoc basis. A system for measuring temperature would be inadequate if it used one function to measure temperatures below the freezing point of water, and another to measure temperatures above freezing. This feature is what a homomorphism is supposed to capture.

Explanatoriness has two elements: a (defeasible) demand to explain intuitive judgments, and the requirement to illuminate structural features of the phenomenon being investigated.

⁹ Treanor's proposal to reduce the ordering of knowledge to the ordering of similarity is not strictly speaking a measurement system; rather it reduces the phenomenon of knowing more to the phenomenon of similarity (Treanor 2013, p. 597).

We should be looking to capture our intuitive first-order judgements about who has more knowledge. There are exceptions. Perhaps in some cases we are mistaken about who knows more. More interestingly, some ordinary judgements might express ad hoc measurement systems that we construct on the fly. For example, if we are comparing the knowledge possessed by students in different school years, we might order students based on the inclusion relation, not because we are committed to this measurement system in general, but just because it works well for comparing these students' knowledge. The existence of ad hoc measurement systems oughtn't to be a surprise: comparing bodies of knowledge is complicated and limited creatures like us ought to use whatever workable systems are ready to hand.

We should also aspire to explain some of the structural features of the phenomenon. A system for measuring knowledge ought to have a sister system for measuring *ignorance* that explains why the growth of knowledge is the destruction of ignorance, (Treanor, 2013, p. 578).¹⁰ A system for measuring knowledge ought also to entail a plausible measurement system for learning. We might also think that a measurement system should explain the convergence of knowledge about subject matters: the idea being that as agents know more about some subject, their epistemic states end up being more and more alike.

One complication with the measurement of knowledge is that in English 'knowledge' is a mass noun,¹¹ meaning that we cannot make sense of 'a knowing', or the question 'how many knowledges does she have?'. We cannot simply count knowledge in English. I don't think that this grammatical point establishes that knowledge is in principle uncountable: many languages use nouns with count readings to refer to knowledge.¹² Rather, it reminds us of the metaphysical point that knowledge is a mental stuff, which must be packaged up to be measured. Water is metaphysically speaking stuff, although in English 'water' does have a count reading: 'can we have two tap waters?'. This means that we cannot simply measure the mass of water; we must measure the mass of spatially or conceptually bounded bodies of water. Similarly, in order to measure the quantity of knowledge, we must measure some bounded body of knowledge. We can think about COUNTING PROPOSITIONS, COUNTING BELIEFS, and COUNTING ANSWERS as having two parts: a criterion for individuating pieces of knowledge (via the object of knowledge, doxastic attitudes, and answers to questions respectively) and a measurement system which maps pieces of knowledge thus individuated onto the structure of the natural numbers.

3 Comparing knowledge

We compare knowledge in various ways in ordinary language. Here is a representative sample:

1 Dorcas knows more than Xenia.

¹⁰ I assume that ignorance is the absence of knowledge, but see Peels (2010).

¹¹ In the sense that it has a mass reading, and no count reading.

¹² Some examples: 'connaissances' (French), 'conoscenze' (Italian), 'conocimientos' (Spanish), 'знания' (Russian). It might even be possible to force a count reading of 'knowledge' in English: think of 'indigenous knowledges', which presumably refers to bodies of knowledge maintained by indigenous groups.

- 2 Laura knows more about Matisse than Ann does.
- 3 Tariq knows a lot more about history than he did.
- 4 Mhairi knows more than Duncan does about who came to the party.
- 5 Rufaida knows more about Ethics than she does about Linguistics.
- 6 Li knows more about why objects fall downwards than the Ancient Greeks did.

These comparisons can be interpersonal (1, 2, 4, 6) or intrapersonal (3, 5), can compare individuals (1, 2, 3, 4, 5), or groups¹³ (6), and can be synchronic (1, 2, 4, 5) or diachronic (3, 6). It is crucial to notice that comparisons of knowledge can either be unqualified (1) or qualified by reference to subject-matters (2, 3, 5), or questions (4, 6). These comparisons may either compare two epistemic states against one issue (2, 3, 4, 6) or against different issues (5). For the next couple of sections, we will focus on question-relative comparisons, returning to subject-matters in section 4.

I want to start with an observation about unqualified comparisons of knowledge like 1. Not to put too fine a point on it, but I think that they are fishy. Shorn of conversational context, claims like ‘Dorcas knows more than Xenia’ seem incomplete, and call for the response ‘what about?’. This is not to say that these sentences are false, meaningless or ungrammatical, but that they are often incomplete in a way that qualified comparisons are not.

In many cases what appear to be unqualified comparisons are covertly qualified. Consider the following exchange between A and B:

A1: What’s the best way to train for a marathon?

B1: Um, I dunno.

A2: No problem - I’ll ask Rich.

B2: Why not ask Lizzie? Lizzie knows much more than Rich does.

B’s comparative claim in B2 is not weird in the way that (1) is. This is because B’s comparison is implicitly question-relative: B is not claiming that Lizzie knows more than Rich all things considered, but that Lizzie knows more about the best way to train for a marathon. When we are looking for an issue, the first place to look will be to the questions under discussion (QUDs) (Roberts, 2012) in a conversation.

Treanor suggests that some of our trouble in understanding some comparisons of knowledge stems from underdescription of the subjects’ knowledge (2013, p. 579). We might think that this explains our trouble with unqualified comparisons. I don’t find this plausible. (1)–(6) are all underdescribed, but it is much more difficult to make sense of the non-relative comparison in (1) than the relative comparisons in (2–6). Another potential explanation is that our difficulty in understanding (1) stems from the complexity of the issue: it’s easier to get a grip on who came to the party than it is on, well, everything. This response doesn’t ring true either: (3) and (6) involve complex issues, and don’t seem problematic.

The lesson I want to draw is that we find it much easier to understand qualified comparisons of knowledge than unqualified comparisons. I think that this provides some suggestive (although not decisive) evidence that in thinking about measuring

¹³ Group comparisons raise further complications: they might be comparing common knowledge, the intersection of individual states of knowledge, sums of individual knowledge, or group-level collective knowledge.

knowledge, we should start with qualified comparisons. Many readers may not share the sense that unqualified comparisons are incomplete. Given the difficulties in building a measurement system for knowledge starting with unqualified sentences, I invite these readers to see this paper as an experiment that starts by taking qualified comparisons of knowledge as central. Even if there turns out to be a serviceable non-counting based account of unqualified sentences, this approach will be of interest for thinking about qualified ascriptions.

4 Questions and inquiry

It is an appealing thought that inquiry is an activity which is aimed at questions (Friedman, 2013, 2019, 2020; Hintikka, 1999; Hookway, 2008). Reports of inquiries typically involve interrogative complements. We say things like ‘Tahlia is inquiring into who came to the party’, and ‘NASA is trying to find out how to sustain human life on Mars.’ According to standard treatments in linguistics, these interrogative phrases denote questions, meaning that these attitudes take questions as its object.

Inquiry is the engine that drives the production of knowledge. This suggests that the metric for knowledge ought to be closely connected to our understanding of progress in inquiry. Progress in inquiry is clearly question-relative. As we shall see in section 5, we can think about a question as a set of alternatives and think about progress in inquiry as involving ruling out alternatives. We start inquiry into a question with all alternatives open, as we progress we can rule out alternatives, until we get to the point where we have the one true answer. This suggests the germ of a measurement system for knowledge: complete ignorance about a question consists in an epistemic state which leaves all alternatives open, complete knowledge involves having ruled out all alternatives, and intermediate states of knowledge can be ranked by how much of the question has been answered.

The connection between knowledge and inquiry gives extra impetus to the idea that the interesting phenomenon is the qualified notion of knowing more about a question and suggests that we might use progress in inquiry to measure amounts of knowledge. In the next section, we will develop this idea, using a partition-based account of questions.

5 Questions

To make the idea of measuring knowledge by questions more precise we need to make some commitments about what questions are.¹⁴ I will appeal to a partition-based framework for thinking about questions (Groenendijk & Stokhof, 1984; Roberts, 2012).¹⁵ Although the details of this framework are a little finicky, we need them to fill out COUNTING ANSWERS. This account relies on distinguishing between complete

¹⁴ This section is heavily indebted to Pavese (2017, n.d.).

¹⁵ For a recent survey of different treatments of questions, including inferential erotetic logic, and inquisitive semantics, see the papers in Cordes (2021).

and partial answers, and giving an account of the role of context, both of which can be done nicely in the partition-based framework.

5.1 Questions as partitions

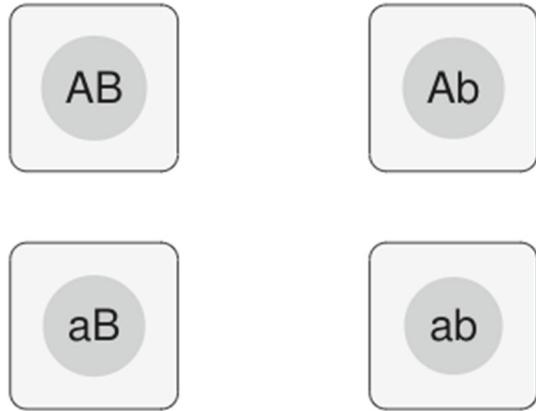
What are questions? Our starting point is that questions are not true or false, but rather have answers which can be assessed for truth. The standard way to explain this is to identify a question not with a proposition, but with a *set* of propositions: those which are its potential answers. The ordinary notion of answer is pretty open: *who was at the party?* might be answered by *John*, by *Thorgy and Sasha came to the party*, by *there wasn't a party* and any number of other propositions. There are various ways to construct a question out of different kinds of answers: Hamblin appeals to the complete set of possible answers (Hamblin, 1958, 1973), Karttunen identifies a question with its correct answer (Karttunen, 1977), and interrogative semantics builds questions out of resolution conditions (Ciardelli et al., 2018). We will be using a partition-based approach which identifies a question with its set of complete answers. A complete answer to a question is a maximally informative proposition, which completely resolves the issue raised by a question. At most one complete answer to a question is correct, and all are incompatible, meaning that the set of answers will divide up a portion of logical space where the presuppositions of the question are met (Lahiri, 2002, p. 10; Mastro, 2010) into a partition. These cells will be mutually exclusive, and the actual world will be contained by (at most) one cell.

Here's a recipe for getting from an interrogative sentence involving a single wh-phrase to the relevant partition. Consider 'who came to the party?'.¹⁶ First, we split the interrogative phrase into two parts: the wh-word ('who'), and the rest of the interrogative ('came to the party'), which we call the question-abstract. We treat the question word as a variable that triggers a domain restriction—'who' to people, 'where' to places, 'why' to reasons and so on—and the question-abstract as expressing a predicate—in this case *came to the party*. Taking this restricted domain, we derive the partition by considering ways to assign all elements in the domain to either the extension or the anti-extension of the predicate. Each assignment is a complete answer: if our domain is just *a* and *b*, then *a came to the party*, and *b didn't come to the party* would be one of the complete answers. The partition is built out of the set of all the complete answers.¹⁷

¹⁶ This is a simplified recipe that will not carry across to whether-interrogatives, multiple wh-interrogatives, or if-questions.

¹⁷ Modes of presentation must play some role in our account of questions. *Whether Ms Marvel came to the party?* and *whether Kamala Khan came to the party?* seem to be distinct questions: one can know, inquire into, or be curious about one of these question but not the other. It is not obvious whether modes of presentation relate to the individuals in the domain of the question, or the entire partition associated with the question. On the individual approach, Ms Marvel and Kamala Khan could both show up in the domain of *who came to the party?* meaning that *Ms Marvel came* and *Kamala Khan came* would could as distinct answers. On the partition approach, the whole partition approach, each individual could only show up once in a question, and *who came to the party?* would need to be disambiguated into two questions associated with each mode of presentation *who came?*_{MM} and *who came?*_{KK}. (the difference being that *who came?*_{MM} would have *whether Ms Marvel came?* as a subquestion and not have *whether Kamala Khan came?* as a subquestion and vice versa). This issue has consequences for the question-relative account of

Fig. 1 Who came to the party?



With our domain restriction in place, *who came to the party?* corresponds to the set {no-one came; only a came; only b came; only a and b came}, giving us the following partition (where X stands for *x came to the party*, and x stands for *x didn't come to the party*) (Fig. 1).

One of the major selling points of the partition-based approach is its ability to distinguish different kinds of answers. Borrowing terminology from (Szabó, 2017), we can distinguish four kinds of answer:

- *Positive minimal answers*, which assign one element in the domain to the extension of the predicate;
- *Negative minimal answers*, which assign one element in the domain to the anti-extension of the predicate;
- *Partial answers*, consisting of combinations of minimal answers with logical connectives;
- *Complete answers*, consisting of complete assignments of all elements in the domain to either the extension or anti-extension of the predicate.¹⁸

Although *who came to the party?* only has four complete answers, it has more minimal answers (such as *a came* and *a didn't come*), and partial answers (such as *either a or b came* and *if a came then b didn't*). Each of these kinds of answer will allow us to rule out some of the cells in the partition. Whereas a complete answer will rule out all but one cell in the partition, minimal and partial answers rule out just some

Footnote 17 continued

measuring knowledge. Consider someone who at t1 knew that Kamala Khan came to the party, and at t2 learns that Ms Marvel came. On the individual approach to modes of presentation, at t2 A knows more about who came to the party than they did at t1, because *whether KK came?* and *whether MM came?* are both positive minimal answers to *who came?*, and they now know both answers. On the partition approach, A still gains more knowledge by t2, but she now knows answers to two questions: *who came?_{KK}* and *who came?_{MM}*. See Aloni (2008) for a proposal which relativises questions to conceptual covers (a version of the partition approach), Aloni and Jacinto (2014) for an application of this approach to knowing who, and Stanley (2011, pp. 106–109) for the suggestion that some of the context-sensitivity of knowledge-wh has close relation to Frege puzzles.

¹⁸ As defined, partial answers will include complete answers, although for the most part we will be interested in non-complete partial answers.

of the cells in the partition. For example, *a came* rules out *aB* and *ab*, and *either a or b came* rules out *ab*. We can often reach a complete answer by combining different partial answers, in the simplest case by combining all of the minimal answers.

The important point that there is a natural way of ranking answers by considering how many cells of a question they rule out. *If a came then b came* only rules out one cell (*Ab*), *a came* rules out two cells (*aB* and *ab*), and *a and b came* rules out three cells (*aB*, *Ab*, and *ab*). We can think of how many cells are ruled out as a metric on how far a partial answer gets us toward successfully answering the question: some answers help a little bit, and others help a lot.

We can use this metric to understand progress in inquiry: the further along in inquiry one is, the more complete answers to the question are ruled out. Consider a climate scientist who is investigating the average temperature in the holocene: they will start off with some large range of possible temperatures (perhaps running from the highest known average to the lowest known averages) along with a level of granularity of answers that reflects their practical interests and the quality of evidence, and as they build up evidence they will make progress by ruling out possibilities which are incompatible with the evidence. Some pieces of evidence will make big contributions, ruling out a lot of possibilities, and others will make smaller contributions, ruling out fewer possibilities. The suggestion of COUNTING ANSWERS is that we can use this metric—the number of complete answers ruled out—to measure an agent's knowledge.¹⁹

This theory of questions is a close fit with the current best theory of subject-matters, which treats subject-matters as partitions. There has been recent interest in subject-matters (Yablo, 2014, 2017, n.d.; Yalcin, 2018), taking inspiration from (Lewis, 1988a, 1988b). According to Lewis' approach, a subject-matter is an equivalence relation on worlds, corresponding to a partition where each cell is a set of worlds that is similar with respect to the relevant part of the world.²⁰ For example, the subject-matter *the 1990s* groups worlds into mutually exclusive sets of worlds which are alike with respect to what happened in the 1990s. We don't have anything quite as elegant as the recipe above for reaching the partition corresponding to a subject-matter, but we can think of subject-matters as corresponding to temporally (*the 1990s*), spatially (*Scotland*), or thematically (*Birds*) circumscribed parts of the world, on top of which the equivalence relation is constructed.

The close relation between subject-matters and questions is reflected in the ease of glossing subject-matters using interrogatives.²¹ For example, the subject matter *the 1990s* can be glossed as *what happened in the 1990s?*, where each cell in this issue will correspond to a complete answer specifying a complete timeline for events from the 1990s. Each answer to this question will be a set of worlds that are similar with

¹⁹ Related ideas play a role in the treatment of sentences like 'Jane knows *in part* who came to the party'. See Pavese (2017, n.d.), Groenendijk and Stokhof (1994), Williams (2000), Beck and Sharvitt (2002), and Lahiri (2002).

²⁰ Lewis works with exact similarity, but we might weaken this relation to relevant similarity, while holding onto the idea that relevance divides up worlds into mutually incompatible possibilities (see below).

²¹ The goal of this comparison is not to reduce subject-matters to questions, but to point out the fact that we intuitively shift between subject-matters and questions, which is neatly explained if both are partitions of possible worlds.

respect to a temporally bounded portion of the world. The subject-matter *Scotland* can be glossed as *what's up with Scotland?*. Each cell in this partition will correspond to a complete description of the geography, history, politics, culture (and so on) of Scotland. This partition is a set of sets of worlds that are similar with respect to a spatially bounded portion of the world, corresponding (at first pass) to the current borders of Scotland. *What's up with Scotland?* is unwieldy (and has unfortunate pragmatics), so it might be easier to think about *Scotland* as being equivalent to a long conjunctive question from which sub-issues might be detached; something like *what is the geography of Scotland, and what is the history of Scotland, and what is the culture of Scotland [...] ?*. The subject-matter *Birds* can be glossed as *what's up with Birds?* and this issue will correspond to a partition in which each cell is a full description of which birds there are, how each type of bird behaves, what their characteristic male and female coloration is, and so on. Here the equivalence relation is constructed on top of a biological portion of the world—the creatures in the Avialae group—with each cell consisting in a maximally informative way that bird-related portion of the world might be.

Although this Lewisian sketch is far from a complete treatment of subject-matters, it opens up some useful doors, allowing us to think about knowledge of a subject-matter as analogous to knowledge of the answer to a question: it consists in knowledge which narrows down a range of possibilities, with the limit case being knowledge that rules out all but one alternative (Yablo, n.d.). Each of our three examples involves a rather complicated subject-matter, with a large number of cells, meaning that we cannot neatly list the alternatives in the partition.²² Although this lack of detail might seem like an explanatory failure, I think it reflects the limits of a general characterisation of subject-matters. A detailed unpacking of the partition corresponding to *Birds* would effectively be a summary of the topics of an expert Ornithologist's knowledge, plus the topics of current investigation, and those topics within their field they are yet to investigate. Saying that *Birds* is a set of ways the bird-related portion of the world might be gives us a pretty thin characterisation of this topic, but that doesn't mean it's not a useful tool for thinking about it as a subject of knowledge, understanding, and a background for comparative knowledge claims.

5.2 Context and questions

To fill out our picture of questions, we need to address the role of context.²³ In section 3 we saw that context can provide a question when none is articulated by the comparative sentence, but the more complex case is when context affects the meaning of an articulated interrogative. There are five mechanisms to consider: domain restriction, modulation of the question abstract, selecting a kind of answer, question substitution, and modes of identification.

²² Not all subject-matters are so complicated: *The Party* might correspond to a partition in which each cell is a list of possible party-goers, notable events, drinks and food served, and times at which people left.

²³ For an invariantist account of interrogatives, see Braun (2006, 2011), and for assessment, see Aloni and Jacinto (2014).

First, context might supply a domain restriction, meaning that the question only considers the application of the predicate to a salient set of individuals. If I utter ‘who came to the party?’ in a conversation about what our friends did at the weekend, the domain will be our friends.

Secondly, context can affect the meaning of the question-abstract. The default meaning for ‘came to the party’ might be *was at the party for any amount of time*, but if we want to know whether people have alibis for the whole evening, ‘came to the party’ might mean *was at the party for the whole night*.

Thirdly, context can affect what kind of answer we are interested in knowledge of. In many cases, we will treat a partial answer as giving us enough information about an issue to settle it for us, giving us what linguists call a *mention-some* reading. If I am interested in skinning a cat, just one method for skinning a cat will resolve my curiosity, even though a complete answer to *how to skin a cat?* would be an assignment of *is/isn’t a way to skin a cat* to every salient method. To account for this, we might introduce the notion of a *resolving* answer, which refers to answers that do enough to meet inquirers’ needs. Resolvingness will not play a role in our measurement system, since having more information than is required for a resolving answer counts as knowing more about the question.

Fourthly, questions often are associated with methods for identifying the objects in the domain (Aloni, 2002). Consider the question *who is Chris Whitty?*. In lots of ordinary contexts, linguistic identification via salient properties will suffice for an answer; so that knowing that Chris Whitty is the UK government’s chief medical advisor will suffice for knowing who Chris Whitty is. In other contexts, different modes of identification will be important. In a line-up situation, an answer will require visual or demonstrative identification that distinguishes Whitty from the other people in the line-up.

Finally, context can switch the question under consideration. Consider:

- 7 Donna knows more than Ricky about whether the standard model of physics is correct.

This question only has two possible answers (*the standard model is correct* and *the standard model is not correct*), meaning that unless one person knows the answer, it is not possible to compare states of knowledge by counting complete answers. In order to get a comparison, we need to substitute the whether question for something more fine-grained, such as *what is the best evidence about the correctness of the standard model?*. This question has a lot of answers, so provides a better background for the comparison of epistemic states. Polar questions are the simplest case of question substitution, but I suspect that it will be a common phenomenon.²⁴ Question-substitution in interrogative clauses is a more substantial and less systematic phenomenon than the other kinds of context-sensitivity, and it would be an interesting project to try to understand it better.

There may be cases in which context seems to fail provide sufficient material to yield a determinate question. An unqualified comparison may be made in a conversation

²⁴ Question substitution can help us explain cases in which some parts of a subject seem to matter more than others. When a physicist asks how much someone knows about physics, the research and teaching environment—and perhaps personal quirks—will generate an implicit list of questions or topics within physics.

in which there is no QUD, or a qualified comparison may be made in a conversation where the participants have different views about the salient domain. Exactly what one wants to say about the content of sentences and speech acts which are made in this kind of situation will depend on one's general view about context-sensitivity, but my suspicion is that all views will end up saying that this kind of comparison are in some sense incomplete.²⁵ This incompleteness means that in some cases comparison will require a degree of discourse repair, in the form of substituting a plausible question into the context to be the background for the comparison.

This discussion of the role of context in questions points towards a couple of different ways in which conversational context might also affect which partition is expressed by a subject-matter denoting phrase.

First, context might determine the granularity of the partition. According to Lewis, subject-matters group worlds by exact similarity, meaning that a cell in *the 1990s* will be a set of worlds which are atom-for-atom replicas of one another with respect to the period from the 1st January 1990 to the 31st December 1999. If we don't care about the 90s to this meticulous degree, we might instead be interested in a subject-matter where the worlds are relevantly similar up to some degree of granularity. For most purposes, it might be enough to group together worlds in which the same historical events happened, without paying attention to what John Major had for breakfast on the 4th February 1991, or whether the 1997 UK election result was declared at 4.00am or 4.01am.

Secondly, context can render some parts of an issue more salient than others. In a context of a conversation about the history of different Enlightenment periods, the subject matter denoted by 'Scotland' (as in 'John knows about Scotland') is really *The Scottish Enlightenment*, and comparative claims about who knows more about Scotland should be answerable to this historical period, ignoring whisky varieties, the highland clearances, and the archaeology of Orkney. In some cases, it might be clear that there is some restriction intended, but not what that restriction is, meaning that we will similarly need to engage in discourse repair to provide a subject-matter as the background for comparative knowledge claims.

6 Counting answers

With the partition-based theory of questions in place, we can develop the idea that knowing more about a question is a matter of knowing more of the answer to that question. Counting complete answers ruled out provides us with plausible ways of measuring progress in inquiry, and knowledge about a question. Returning to the idea that 'knowledge' is a mass noun, we can put the point by saying that the division of logical space provided by a question provides us with a way to parcel up knowledge into countable pieces. This gets us to COUNTING ANSWERS:

²⁵ Grindrod and Borg (2019) argues that many appeals to QUDs to resolve context sensitivity raise similar problems.

COUNTING ANSWERS: A knows more than B about a question Q, just in case A's knowledge allows her to rule out more of the complete answers in Q than B's knowledge does.

A couple of clarificatory points.

Like COUNTING PROPOSITIONS, COUNTING ANSWERS is an *object-based* measurement system, meaning that we can employ it without making any very substantive commitments about the nature of the knowledge relation (see footnote 6 for some complications).

Ruling out refers to the way one's knowledge can close off certain possibilities, making them irrelevant for the purposes of practical and theoretical reasoning. I will remain neutral on the question of how much evidence is required for ruling out.

This view doesn't measure knowledge by counting just any kind of answer ruled out: we are interested only in *complete* answers to a question.²⁶ We wouldn't get very far counting minimal answers, or partial answers.²⁷ If we were to count minimal answers, we would lose the contribution of logically complex knowledge (since a conditional proposition like *if a came then b didn't* does not entail any minimal answers), and if we were to count partial answers we would pretty quickly get into worries about counting infinities (since partial answers can involve disjunction, we could use disjunction introduction to turn any set of partial answers into an infinite set).

Another feature of the account is that it puts positive and negative minimal answers on par. If the question is *who came to the party?* our domain is a, b, and c and a and b came and c didn't, then if X knows that a came, and Y knows that c didn't come, then X and Y have the same amount of knowledge. X can rule out four possible complete answers (aBC; aBc; aBC; abc), and so can Y (ABC; aBC; aBc; AbC). This result might seem strange, since X is closer to a having the complete list (*a* and *b* came), but it reflects the fact that X and Y are equally close to the complete answer (*a* and *b* came, and no-one else did).²⁸

This framework measures knowledge. It is silent on how we should measure justification or closeness to the truth. If one agent can rule out one possible complete answer, then she will know more than either (i) an agent who has non-knowledge-level justified beliefs that two possible complete answers are incorrect, or (ii) an agent who has an unjustified true belief that rules out many complete answers. It also doesn't differentiate between knowledge about different questions: knowing a lot about Drag Race contestants and knowing a lot about fundamental physics are both ways to know a lot.

COUNTING ANSWERS allows us to order agents with respect to how much they know about a question, it has a non-arbitrary zero point—complete ignorance about

²⁶ Are all complete answers equal? COUNTING ANSWERS is committed to this, but we might have our doubts: if A knows that the number of planets is either 7, 8, or 9, and B knows that it is either 9, 100, or 972, we might think that A knows more about how many planets are in the solar system. I'm tempted to think that in this kind of case the increased knowledge is concerned with salient sub-questions, such as *whether the number of planets is less than 100?*

²⁷ See Pavese's discussion what she calls the COUNTING APPROACH, which uses Hamblin-answers (Pavese, 2017, n.d.).

²⁸ There might well be situations where we employ an ad hoc measurement system that counts minimal positive answers (see section 7.3.).

some question—and puts numerical values on a mass-like phenomenon. It is tempting to think of it as a ration scale for knowledge. To establish this, we would need to show that each complete answer ruled out corresponds to the same increment in knowledge. This assumption is not obvious. It could be that the increment in knowledge from two open possibilities to one possibility (i.e. fully answering the question) is greater than the increment from every possibility being open to one possibility being ruled out.²⁹

COUNTING ANSWERS can be naturally extended to give a theory of the measurement of ignorance and learning. If knowledge is measured by counting the number of complete answers to a question that are ruled out, we can measure ignorance by counting the number of complete answers left open. This explains why the measurement of knowledge mirrors the measurement of ignorance, and why knowledge is the destruction of ignorance.³⁰ Similarly, we can measure learning relative to a question, by ranking changes in epistemic states by the change in the number of complete answered ruled out.

We have seen that on the most popular theory of subject-matters, subject-matters and questions have closely related structures. Applying similar ideas to subject-matters gives us:

COUNTING ALTERNATIVES: A knows more than B about a subject-matter S, just in case A's knowledge allows her to rule out more alternatives in S than B's knowledge does.

Just as COUNTING ANSWERS requires counting the number of complete answers to a question which are ruled out by an agent's knowledge, COUNTING ALTERNATIVES requires counting the number of relevantly similar groupings of worlds ruled out by an agent's knowledge. Giving an account of subject-matters, and how they interact with context is an interesting issue, but is best left for another day.

This allows us to give the following account of knowing more about an issue:

COUNTING ISSUES: A knows more than B about an issue I, just in case A's knowledge allows her to rule out more of the cells in the partition corresponding to I than B's knowledge does.

6.1 Contextualism, invariantism, and subject-sensitivism

As we have articulated it, COUNTING ISSUES is an account of what it is to know more, not an account of sentences that compare knowledge. Here we face an important choice point in rounding out the account. We might give different accounts of the origins of the background questions, appealing to context, a question which doesn't vary, or something about the agent's position to supply questions to provide the background metric.

²⁹ Richard Pettigrew informs me that mathematics exams in Oxford were historically marked in a non-linear way: the answers to questions were first marked out of ten, then the marks were squared, incentivizing more complete answers to questions.

³⁰ Here we are focusing on first order ignorance. On higher-order ignorance, see Wilholt (2020).

These options have structural similarities to the views in the debate about the semantics of propositional knowledge ascriptions so, following the lead of the debate about knows-wh ascriptions (Braun, 2006, 2011; Parent, 2014) I will adopt the terminology of contextualism, invariantism, and subject-sensitivism. However, there is no straightforward correspondence between—say—contextualism about comparative sentences, and contextualism about propositional knowledge ascriptions. The former claims that there is context-sensitivity in the (articulated or unarticulated) question involved in comparison of knowledge, and the latter claims that the verb ‘knows’ is context-sensitive. Just as it is possible to be a contextualist about the interrogative portion of knowledge-wh ascriptions, and a subject-sensitive invariantist about the verb in knowledge-that ascriptions (see Stanley, 2005, 2011), positions in these two debates are separate. My official position combines invariantism about ‘knows’ with contextualism about comparative sentences and knowledge-wh ascriptions.

I will suggest that the balance of reasons favours contextualism about comparative sentences, but developing the alternatives helps us to see some of the features of contextualism.

According to contextualism, the background issue for a comparison will either come from an articulated interrogative clause (interpreted in line with the context), or from the issues under discussion in the conversation. This theory entails the following semantics for qualified and unqualified comparisons:

- A sentence of the form “A knows more than B about Q/S” is true in a context C iff A’s knowledge rules out more of the cells in the issue I expressed by “Q” or “S” in C than B’s knowledge does.
- A sentence of the form “A knows more than B” is true in a context C iff A’s knowledge rules out more of the cells in the contextually salient issue I than B’s knowledge does.

If the context does not provide suitable material to yield a fully articulated question, then the comparison may be indeterminate. If I say ‘Ann knows more than Bernard’, but it is indeterminate whether we are talking about theology, English sparkling wine, or music, the sentence will fail to express a determinate comparison. In such cases, we may need to engage in discourse repair, as we noted in section 5.

The invariantist claims that all comparisons either directly or indirectly involve a single invariant *ur*-question. Perhaps this question is something general like *what’s up?* or *how are things?* or some metaphysically special question like *what is the correct fundamental theory of reality?* All unqualified comparisons will be made against the *ur*-question, which is the direct case. Overtly and covertly qualified comparisons will be made against the background of the conversationally salient question, but indeterminacy in this question will be resolved by treating it as a subquestion of the *ur*-question, which is the indirect case.

The subject-sensitive account claims that the subject’s epistemic state provides the background for epistemic comparisons, much as the subject’s stakes provide the standards for knowledge according to the corresponding view about propositional knowledge (Stanley, 2005). The background questions might be those an agent can articulate, those she is inquiring into, or those of her theoretical paradigm. This

approach applies most smoothly to intrapersonal comparisons relating to a single inquiry. When we make an unqualified comparison, such as:

8 Arlette knows more than she did

plausibly what we mean is that Arlette can rule out more of the possible complete answers to her questions than she used to be able to. When we make a qualified comparison, like:

3. Tariq knows a lot more about how to do carpentry than he did

we take the question given by the qualifier phrase (*how to do carpentry?*), and filter it through the agent's conceptual schema (taking the subquestions which the agent is capable of articulating), then consider how much of that filtered question the agent can answer.

I will opt for a contextualist treatment of the semantics of natural language comparisons of knowledge. The approach provides the most natural explanation of why comparisons often feel underspecified or incomplete. It also traffics in simpler questions, making it easier to evaluate comparative claims. This approach also nicely separates issues about the importance of questions from the measurement of knowledge: its aim is to give an account of knowing more about questions, irrespective of their importance. Because they compare relative to very large questions, both the invariantist and subject-sensitive approaches seem to require a ranking of the importance of different questions within those big questions.³¹ The Invariantist also faces difficulties with changes of questions. If two agents have different conceptual schema—or an agent expands her schema (Carr, 2015; Pettigrew, 2018a)—there will not be a single question to provide the backdrop to the comparison (see also section 6.2). These considerations are by no means conclusive, and there is work to be done in developing invariantist and subject-sensitive approaches of sentences which measure knowledge.

6.2 Positive features of COUNTING ISSUES

In section 3 I suggested that a good measurement system for knowledge should be general, systematic and should explain our intuitive judgements. The contextualist version of COUNTING ANSWERS does well on these desiderata.

COUNTING ISSUES only gets part way with respect to generality. It naturally explains both qualified comparisons involving questions and subject-matters and unqualified comparisons involving contextually supplied issues. It also applies to both interpersonal and diachronic comparisons. It remains to be seen how the theory can deal with genuinely unqualified comparisons and comparisons between issues, issues we will consider in section 7.

It also does well on systematicity. Predictions about who knows more are generated from a well-understood theoretical framework for thinking about questions and subject-matters that can be applied smoothly to a wide range of domains.

This is not to say that this approach makes measuring knowledge easy. COUNTING ISSUES predicts three barriers to measuring knowledge: lack of clarity about

³¹ This will be especially tricky if the amount that propositions contribute to knowledge varies between worlds, (Levinstein, 2019).

the issue, indeterminacy of issues, and complexity of the states of knowledge being compared. These are general problems that can occur even for good measurement systems. Although the metric system for measuring lengths is a good fit for medium sized dry goods, it is not always easy to measure lengths: some things are hard to hold still (snakes), others have indeterminate lengths (clouds), and others have complicated lengths (springs).

COUNTING ISSUES also explains a wide range of intuitive judgements. In most cases in which we compare states of knowledge, it is pretty clear that there is a question or subject matter in the background, and that working through the details of the agents' knowledge would leave us with a situation in which the agent that knew more could rule out more of the relevant issue. This approach explains why unqualified comparisons seem fishy, and the role of context in altering the meaning of comparisons, and elegantly predicts the relation between knowledge, ignorance and learning in a unified framework.

COUNTING ISSUES also avoids the problems for the counting approaches discussed above.

Whereas there is a very real worry that we make have infinite numbers of beliefs and know infinite numbers of propositions, in the majority of cases, the issues we are interested in have a finite number of complete answers. Although some questions do have infinite numbers of complete answers (an issue we return to in section 7.3), these are special cases, and COUNTING ISSUES has some resources to deal with these cases.

The partition-based account of questions also gives us a clear account of how to individuate the objects which we are counting: they are the smallest cells in the partition corresponding to the issue, that is the answers that output from the semantics for interrogatives and the pragmatic story about the modulation of questions.

Finally, the account offers a clear framework for thinking about how propositions can contribute differentially to knowledge. The examples that Treanor uses to motivate this point are complex, so I'll work with some simple cases. Return to our toy question *who came to the party?*. Consider two agents, one of which knows that *a* came to the party, and the other of which knows that if *a* came to the party, *b* did too. Who knows more about who came to the party? Plausibly, the agent who knows that *a* came to the party knows more. Why is this? *If a came to the party, b did* rules out just one cell (*Ab*), *a came to the party* rules out two cells (*aB* and *ab*), so counting the complete cells ruled out plausibly gets us the right result.

7 Problem cases

Although COUNTING ISSUES clearly deals well with a range of comparisons of knowledge, there are some tricky cases: (i) genuinely non-relative comparisons, (ii) cross-issue comparisons, (iii) issues with an infinite number of cells, and (iv) cases of trivial knowledge.

7.1 Non-relative comparisons

COUNTING ISSUES takes qualified comparisons of knowledge to be the central case, and the strategy with *prima facie* unqualified comparisons like (1) has been to argue that they have an implicit issue-relative structure provided by context, so that:

1. Dorcas knows more than Xenia.

Really means:

1*. Dorcas knows more than Xenia [about I]

One might be dissatisfied by this strategy for a couple of reasons:

- One might think that there are genuinely unqualified comparisons of knowledge, which are such that if we ask ‘what does A know more than B about?’, the response is ‘just in general, A knows more’.
- One might also think that qualified knowledge is just not the interesting phenomenon, that the job of a measurement system for knowledge is to rank people regarding how much they know in general.

Behind both worries lies the concern that COUNTING ISSUES substitutes the difficult and interesting question of who knows more for the easy and uninteresting question of who knows more about some issue.

It is true that COUNTING ISSUES refuses to measure unqualified states of knowledge, but I don’t think that this refusal is as bad as it seems. For starters, COUNTING ISSUES is compatible with measuring knowledge relative to very general questions. Context can provide very general questions. More generally, I don’t see why we should privilege questions about who knows more, all things considered. I suspect that most of the time when we are interested in who knows more, we are interested in domain-specific comparisons. It’s easy to see why you might care about who knows more about plumbing, and difficult to see why you might care about who knows more all things considered.³²

7.2 Comparing across issues

COUNTING ISSUES compares states of knowledge about just one question, meaning that it is not obvious what the approach should say about sentences which compare epistemic states relative to different issues, like (5):

5. Rufaida knows more about Ethics than she does about Linguistics

This sentence is qualified, but each comparison is qualified by a different piece of knowledge, so we cannot simply apply the semantics for COUNTING ISSUES.

The simplest approach would be to count the number of complete answers which the agent can rule out from each question. In cases where the cells have drastically

³² Politicians often make claims about knowing more about specific topics to make unqualified claims about knowing more, which can be a way to use specialized expertise to leverage political power. COUNTING ISSUES gives a nice explanation of why this move is problematic: knowing more about economics does not entail knowing more about the demands of justice.

different numbers of cells, this approach gets into trouble. If issue A has four cells, of which an agent can rule out three, and issue B has sixty-four cells of which she can rule out sixteen, then this approach predicts that she knows more about B than A, although she has maximal knowledge about A.³³

Another possibility would be to construct a *super-issue* out of the two questions and count the number of complete answers which are ruled out. Combining two issues is relatively easy: all we need to do is pick out a portion of logical space reflecting the conjunction of the presuppositions of the two issues, then reflect all the distinctions made by both issues against that portion of logical space. We can then ask how many answers in the super-issue are ruled out.

A final option would be to move from counting the number of cells which an agent can rule out to considering the ruled out cells as a proportion of the cells in the issue. This would mean that ruling out many cells in a complex issue would contribute less to knowledge than ruling out fewer cells in a less complex issue.

Deciding which of these approaches we ought to take is a complex question. I suggest that we accept that COUNTING ISSUES does not deal with sentences like 5, and treat whatever approach we go with as an ad hoc measurement system.

7.3 Infinite answers

In most cases, contextually salient questions will have a finite number of complete answers involving a restricted domain. However, in cases of questions about infinite domains, we will get questions with infinite numbers of complete answers.³⁴ Consider the question *which spacetime points are occupied by an electron?* Each the cells in this partition is a complete assignment of the predicate *is occupied by an electron* to all of the spacetime points. Assuming that there are infinitely many spacetime points, this partition has infinitely many cells. In this case, knowing that any spacetime point is occupied by an electron will rule out an infinite number of cells in this partition, meaning that as soon as two agents both know the location of one electron, they can both rule out an infinite number of cells in the partition.

Much like cross-issue comparisons, issues with infinite numbers of cells resist a simple version of the counting approach. However, there is a work-around that we can use to construct an ad hoc measurement system for these cases: we can count the number of minimal answers which agents know (I assume that we should count both

³³ We might think that this problem can be avoided by counting the number of cells left open, rather than the number of cells ruled out. This approach gets into trouble in similar cases: if A's issue has four cells, of which she has only ruled out one, and B's issue has 64, of which she has ruled out 37, we might think that B knows more about her issues than A does about hers, although B's knowledge leaves more cells open than A's does.

³⁴ One might worry that there will be many cases in which questions involve infinite answers, because all numerical questions will range over the domain of the natural numbers. I suspect that context will typically restrict things. A climate scientist investigating the average temperature in the Holocene might seem to be inquiring into a question with infinite answers, but her context will plausibly provide her both with a temperature range (say -50C to 50C), and a level of granularity of answers corresponding to her interests and the accuracy of her models (say two decimal places), yielding a finite domain. Questions which are about numbers—such as *which numbers are prime?*—pose an important problem for a partition-based understanding of questions, since any answer to this question will be necessarily true or false, and thus identical in all possible worlds. How to address this issue is an important question for theorists of questions.

positive and negative minimal answers, but one might also want to count only positive answers). Counting these minimal answers gets us the intuitively correct result that knowing the locations of more electrons means that one knows more about which spacetime points are occupied by electrons. This approach is not perfect: it does not measure the contribution of partial answers that are not minimal answers, and it will break down if we compare two agents who know the locations of infinite numbers of electrons, but these are going to be hard cases for any theory, and this does give us some tools to deal with them.

7.4 Trivial knowledge-wh

A final worry concerns knowledge of trivial answers (Aloni, 2002). For almost any question, there will be trivial answers which suffice to rule out all but the complete correct answer. For example:

9. Who came to the party?
10. The people who came to the party came to the party.

On the face of it, knowing the proposition expressed by 10 should not suffice for knowing who came. The problem isn't just the use of trivial descriptions—it would be easy to introduce names that allow trivial answers (say if we introduce 'Nigel' as a collective noun for the partygoers). This problem generalises beyond knowing who: *because of the reasons they do* will usually not be an answer to the question *why do objects fall downwards?*. Here the contextualist flavour of the account comes into play again: in the overwhelming majority of situations, context will require some more demanding mode of identification than reference by trivial description: for example, reference by proper names, or by visual identification (see footnote 16).

8 Conclusion

In this paper, I've developed an issue-relative and contextualist approach to measuring knowledge. This approach fits well with our intuitive judgements about comparisons of knowledge, gives a general and systematic picture of the measurement of knowledge, receives powerful support from the connection between knowledge and inquiry, and offers a precise way to understand comparisons of knowledge, ignorance and learning. I cannot hope to have resolved all issues about the measurement of knowledge, but I hope to have made the case that this framework is a useful one, and deserves further development.

In the interests of packaging up our ignorance, here are some outstanding questions:

- How does COUNTING ISSUES compare to other measurement systems for knowledge?
- How does the measurement of knowledge relate to the measurement of *justification*, *truth* and *accuracy*? (Dunn, 2015; Joyce, 1998; Pettigrew, 2016, 2018b; Tang, 2016)
- How might credal states contribute to knowledge? How might graded attitude (Konek, 2016) and graded content (Moss, 2018) approaches be integrated into a question-based system?

- How should we think about the epistemic importance of different questions?
- Does the plausibility of COUNTING ISSUES lend support to issue-relative accounts of belief and knowledge?
- Can we give a question-based measurement system for truth or content? (see Yablo, 2014; Treanor, 2018).

Acknowledgements Thanks to Ludovica Adamo, Haixin Dang, Bruno Jacinto, Jason Konek, Sophie Archer, Alexander Dietz, Mary Edwards, Simon Graf, Richard Gray, Patrick Greenough, Elizabeth Irvine, Heather Logue, Orestis Palermos, Andrew Peet, Richard Pettigrew, Justin Snedegar, Mona Simion, Jonathan Webber, and audiences at the Cardiff Departmental seminar, and the EJK seminar and the 2017 postgraduate reading party in St Andrews. Special thanks to two reviewers for *Synthese*. The project leading to this publication has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant Agreement No. 818633).

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Aloni, M. (2002). Questions under cover. In D. Barker-Plummer, D. Beaver, J. van Benthem, & P. Scotto de Luzio (Eds.), *Words, proofs, and diagrams*. CSLI.
- Aloni, M. (2008). Concealed questions under cover. *Grazer Philosophische Studien*, 77(1), 191–216.
- Aloni, M., & Jacinto, B. (2014). Knowing Who: How perspectives and contexts interact. In F. Lihoreau, & M. Rebuschi (Eds.), *Epistemology, context and formalism*. Synthese Library.
- Beck, S., & Sharvitt, Y. (2002). Pluralities of questions. *Journal of Semantics*, 19, 105–157.
- Bird, A. (2007). What Is scientific progress? *Nous*, 41(1), 64–89.
- Braun, D. (2006). Now you know who Hong oak yun is. *Philosophical Issues*, 16(1), 24–42.
- Braun, D. (2011). Knowing how and knowing answers. In J. Bengson & M. Moffett (Eds.), *Knowing how: Essays on knowledge, mind, and action I* (pp. 244–260). Oxford University Press.
- Carr, J. (2015). Epistemic expansions. *Res Philosophica*, 92(2), 217–236.
- Ciardelli, I., Groenendijk, J., & Roelofsen, F. (2018). *Inquisitive semantics*. OUP.
- Cohen, S. (1986). Knowledge and context. *Journal of Philosophy*, 83(10), 574–583.
- Cordes, M. (2021). *Asking and answering: Rivalling approaches to interrogative methods*. Narr Francke Attempto Verlag.
- DeRose, K. (2009). *The case for contextualism: Knowledge, skepticism, and context* (Vol. 1). Oxford University Press.
- Dunn, J. (2015). Reliability for degrees of belief. *Philosophical Studies*, 172(7), 1929–1952.
- Friedman, J. (2013). Question-directed attitudes. *Philosophical Perspectives*, 27(1), 145–174.
- Friedman, J. (2019). Inquiry and belief. *Nous*, 53(2), 296–315.
- Friedman, J. (2020). The epistemic and the zetetic. *Philosophical Review*, 129(4), 501–536.
- Goldman, A. (2001). Experts: Which ones should you trust. *Philosophy and Phenomenological Research*, 63(1), 86–110.
- Grindrod, J., & Borg, E. (2019). Questions under discussion and the semantics/pragmatics divide. *Philosophical Quarterly*, 69(275), 418–426.
- Groenendijk, J., & Stokhof, M. (1984). *Studies on the semantics of questions and the pragmatics of answers*. Amsterdam, The Netherlands. <http://dare.uva.nl/record/123669>
- Groenendijk, J., & Stokhof, M. (1994). Interrogatives and adverbs of quantification. In K. Bimbo & A. Mate (Eds.), *Proceedings of the 4th symposium of logic and language*. Budapest.

- Habgood-Coote, J. (2018). Knowledge-how, abilities, and questions. *Australasian Journal of Philosophy*, 97(1), 86–104.
- Hamblin, C. (1958). Questions. *Australasian Journal of Philosophy*, 36(3), 159–168.
- Hamblin, C. (1973). Questions in Montague English. *Foundations of Language*, 10, 41–53.
- Hintikka, J. (1999). *Inquiry as inquiry: A logic of scientific discovery*. Kluwer Academic Publishers.
- Hookway, C. (2008). Questions, epistemology and inquiries. *Grazer Philosophische Studien*, 77(1), 1–21.
- Joyce, J. M. (1998). A nonpragmatic vindication of probabilism. *Philosophy of Science*, 65(4), 575–603.
- Karttunen, L. (1977). Syntax and semantics of questions. *Linguistics and Philosophy*, 1, 3–44.
- Konek, J. (2016). Probabilistic knowledge and cognitive ability. *Philosophical Review*, 125(4), 509–586.
- Lahiri, U. (2002). *Questions and answers in embedded contexts*. OUP.
- Levinstein, B. (2019). An objection of varying importance to epistemic utility theory. *Philosophical Studies*, 176(11), 2919–2931.
- Lewis, D. (1988a). Relevant implication. *Theoria*, 54(3), 161–174.
- Lewis, D. (1988b). Statements partly about observation. *Philosophical Papers*, 17(1), 1–31.
- Lewis, D. K. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, 74(4), 549–567.
- Masto, M. (2010). Questions, answers, and knowledge-wh. *Philosophical Studies*, 147(3), 395–413.
- Moss, S. (2018). *Probabilistic knowledge*. OUP.
- Parent, T. (2014). Knowing-wh and embedded questions. *Philosophy Compass*, 9(2), 81–95.
- Pavese, C. (n.d.). Questions and the structure of answers, 1–46.
- Pavese, C. (2017). Know how and gradability. *Philosophical Review*, 123(3), 345.
- Peels, R. (2010). What is ignorance? *Philosophia*, 38(1), 57–67.
- Pettigrew, R. (2016). *Accuracy and the laws of credence*. Oxford University Press.
- Pettigrew, R. (2018a). The population ethics of belief: In search of an epistemic theory X. *Nous*, 52(2), 336–372.
- Pettigrew, R. (2018b). What is justified credence? *Episteme*, 18(1), 16–30.
- Roberts, C. (2012). Information structure in discourse: Towards an integrated formal theory of pragmatics. *Semantics and Pragmatics*, 5(6), 1–69.
- Roelofsen, F. (2019). Semantic theories of questions. In *Oxford research encyclopedia of linguistics*. OUP.
- Schaffer, J. (2005). Contrastive knowledge. In T. S. Gendler & J. Hawthorne (Eds.), *Oxford studies in epistemology 1*. Oxford University Press.
- Schaffer, J. (2008). Knowledge in the image of assertion. *Philosophical Issues*, 18, 7–22.
- Stanley, J. (2005). *Knowledge and practical interests*. Oxford University Press.
- Stanley, J. (2011). *Know how*. Oxford University Press.
- Szabó, Z. G. (2017). Finding the question. *Philosophical Studies*, 174(3), 779–786.
- Tang, W. H. (2016). Reliability theories of justified credence. *Mind*, 125(497), 63–94.
- Treanor, N. (2013). The measure of knowledge. *Nous*, 47(3), 577–601.
- Treanor, N. (2014). Trivial truths and the aim of inquiry. *Philosophy and Phenomenological Research*, 89(3), 552–559.
- Treanor, N. (2018). Truth and epistemic value. *European Journal of Philosophy*, 26(3), 1057–1068.
- Treanor, N. (2019). The proper work of the intellect. *Journal of the American Philosophical Association*, 5(1), 22–40.
- Wilholt, T. (2020). On knowing what one does not know: Ignorance and the aims of research. In J. Kourany & M. Carrier (Eds.), *Science and the production of ignorance* (pp. 195–218). MIT Press.
- Williams, A. (2000). Adverbial quantification over (interrogative) complements. In *WCCFL 19 proceedings*, (pp. 574–587).
- Wisniewski, A. (1995). *The posing of questions: The logical foundations of erotetic inferences*. Springer.
- Yablo, S. (n.d.). Knowing about things. Retrieved February 6, 2017, from http://www.mit.edu/~yablo/home/Papers_files/KnowingAboutThings.pdf
- Yablo, S. (2014). *Aboutness*. Princeton University Press.
- Yablo, S. (2017). Parts and differences. *Philosophical Studies*, 173(1), 141–157.
- Yalcin, S. (2018). Belief as question-sensitive. *Philosophy and Phenomenological Research*, 97(1), 1–32.