

Fragmentation and information access*

Adam Elga and Agustín Rayo[†]

March 1, 2015

Abstract

In order to predict and explain behavior, one cannot specify the mental state of an agent merely by saying what information she possesses. Instead one must specify what information is available to an agent relative to various purposes. Specifying mental states in this way allows one to accommodate cases of imperfect recall, represent the mental states of confused or fragmented subjects, and shed light on the difference between propositional knowledge and know-how.

Contents

1	Introduction	2
2	Limitations on information access	2
2.1	The motivating idea	2
2.2	Imperfect recall: Jack’s memory	3
2.3	Aha!-moments: a word puzzle and a metapuzzle	4
3	Access tables and indexed mental states	5
4	Is indexing necessary?	7
5	Further applications	10
5.1	Mental fragmentation	10
5.2	How to characterize fragmented mental states	12
5.3	Indexing and abilities	13
5.4	Knowledge-how versus knowledge-that	15

*Thanks to Ross Cameron, David Chalmers, Jonathan Cohen, Keith DeRose, Cian Dorr, Hartry Field, Branden Fitelson, David Hunter, Frank Jackson, Shivaram Lingamneni, Michael Rescorla, Robert Stalnaker, Jason Stanley, Bruno Whittle, audiences at the 2008 Arizona Ontology Conference, Brown University, the Catholic University of Peru, CUNY, the National Autonomous University of Mexico, the University of Bologna, UC Berkeley, UC Riverside, UC Santa Cruz, University of Leeds, University of Paris (IHPST), University of Oslo, University of Texas at Austin, and Yale University, the Corridor reading group (on two occasions), and participants in a graduate seminar session at Rutgers University and a Fall 2011 joint MIT/Princeton graduate seminar. The initial direction of this paper was enormously influenced by conversations with Andy Egan.

[†]ae@adamelga.net, arayo@mit.edu

1 Introduction

A soldier walks an intricate path to successfully evade the mines littering a mine-field.¹ A student scratches out just the squiggles on a piece of paper that earn her top marks. A bank manager turns the dials of a safe in the precise way required to open it. What helps explain each success is that the subject has access to appropriate information.

The aim of this paper is to draw attention to cases involving limited information access. Our main claim is that in order for an information-based² theory of behavior to accommodate such cases, it must be “indexed”: it must determine what information is available to an agent *relative to various purposes*.³

2 Limitations on information access

2.1 The motivating idea

Sometimes different bits of information are available to an agent for different purposes. To see why, cast your mind back to the early 1980s, well before telephone numbers were searchable online.

In those days—the early days of computer hacking—“reverse phone books” were prized by hackers. A reverse phone book is just a phone book whose entries are sorted by telephone number instead of by name. Such books were valuable because they enabled one to easily figure out the name associated with any given phone number. A hacker would sometimes need to swindle information from the person at a given phone number. He would use the reverse phone book to get the name of that person, and then call the number and use his knowledge of the name to gain the person’s confidence.

Compare a phone book for a neighborhood with its corresponding reverse phone book. The hacker doesn’t care about the phone book at all, but desperately wants the reverse phone book. Why? Not because the reverse

¹Cf. Braddon-Mitchell (2007).

²For a theory of behavior to be information-based is for it to deploy intentional (or representational) notions to describe a subject’s mental state. In contrast, a non-information based theory might seek to account for a subject’s behavior purely in physical or non-intentional psychological terms.

³Theories congenial to indexing include Brown and Priest (2004), Egan (2008), Lewis (1982), Parikh (2008).

phone book contains different information than the ordinary one. After all, the books contain exactly the same listings, just in a different order. Rather, the hacker prizes the reverse phone book because of the way it allows him to access that information.⁴

A hacker equipped with a reverse phone book is able to quickly answer questions of the form: "Who has phone number 545-3899?". In contrast, answering such a question using just an ordinary phone book would require painstakingly examining every listing on hundreds of pages.

Let us use "accessible" to mean "immediately available". Then we can summarize the discussion so far as follows: information can be represented in a way that makes it accessible for some purposes, but inaccessible for others (Stalnaker 1991).⁵

2.2 *Imperfect recall: Jack's memory*

The above point applies to human memories as much as it does to phone books. Often our memories provide fast access to some information for one purpose but not another. For example:

Jack's memory Jack has a neighbor he sees only infrequently. The neighbor's name is "Beatrice Ogden", and she lives in apartment 23-H. If asked "What is the name of the person in 23-H?" Jack is disposed to groan, scratch his head, mutter "I know this, don't tell me..." but be unable to answer. But if instead asked "How do you know Beatrice Ogden?", Jack is disposed to immediately reply, "She's the person in 23-H."

Jack possesses the information that Beatrice Ogden is the person in 23-H, but is able to access it for some purposes and not for others. He enjoys the sort of access that would be helpful when attempting to direct a letter addressed "To: Beatrice Ogden", but not the sort of access that would be helpful when going over to 23-H and calling a greeting through the door.

⁴Talk of "information" is in one respect misleading. Saying that a subject has some information at least suggests that the information is *true*. But the truth value of the "information" will play no role in our discussions. So we will use "information" non-factively, so that false information counts as information, too.

⁵Indeed, as the designers of computer databases are well aware, this is not the exception but the rule. Often setting up a database to make its information easily accessible for one purpose will make that information less accessible for other purposes.

2.3 *Aha!-moments: a word puzzle and a metapuzzle*

The difference between a batch of information being present in a system and its being accessible for a certain purpose can be used to represent certain kinds of Aha!-moments in reasoning. As an example, consider the following word puzzle:

Is there a word of English ending in the letters “MT”?⁶

Think about this puzzle for a moment before turning the page.

⁶This puzzle, and the use of it in the present context, is adapted from a very similar puzzle in Powers (1978, 340–341). See also Crimmins (1992).

DREAMT.

Now for the metapuzzle. Even before hearing the puzzle you knew perfectly well that “dreamt” is a word of English, and you knew how to spell it. Indeed, you knew that “dreamt” is a word of English containing the letters “MT”. (If before hearing the puzzle someone had excitedly announced to you, “ ‘Dreamt’ is an English word ending in the letters ‘MT’ !” you might have thought to yourself, “I knew *that* already.”)⁷

This means that information you already possessed—that “dreamt” is a word of English ending in “MT”—is sufficient to solve the puzzle. The metapuzzle is: given that you had in advance information that suffices to solve the puzzle, why wasn’t the puzzle trivial to solve?

We suggest that when you solve the puzzle, information you possessed all along—the information that “dreamt” is a word of English ending in “MT”—becomes accessible for a new purpose. It becomes accessible for the purpose of answering the question “Is there a word of English ending in the letters ‘MT’ ?”^{8,9}

Lesson: sometimes an Aha!-moment—such as the mental change that occurs when one first solves the “dreamt” puzzle—crucially involves a change in what information is accessible to a subject for a particular purpose.

3 Access tables and indexed mental states

We have claimed that to accommodate cases involving limited recall, we must represent mental states using “indexing”: we must specify what information is available to an agent relative to various purposes. It is time to give a concrete example of how such indexing might work.

Start with some terminology. Say that an “elicitation condition” for an agent is a choice situation for that agent. For example in the case of Jack, one elicitation condition is being asked “who is in apartment 23-H?”. Another one is being asked “who is Beatrice Ogden?”. Elicitation conditions often prompt an agent to deploy information for a particular purpose, so it is sometimes convenient to think of elicitation conditions as purposes.

Now for the example. Let us associate with each mental state an “access table”—a table specifying what information is available relative to each

⁷This case and its use in the above argument is drawn directly from (Powers 1978, 341).

⁸Philosophy of mind mavens may find this reminiscent of the “ability” reply to the Mary argument. That connection is pursued in Rayo (2013, ch. 4). See also Cath (2009).

⁹Some will say that in addition to existing information becoming accessible to you for new purposes, something else happened as well: you acquired *new* information. We doubt this (Stalnaker 1984, 1991), but won’t argue the point here. See Elga and Rayo (2014).

elicitation condition. For example, below is an access table for the Phonebook Creature. Like all of the access tables in this paper, this one is abridged and simplified for brevity:

<i>Elicitation Condition</i>	<i>Information accessible relative to that condition</i>
Has recently been asked only the question “What is Selma K. Petersen’s phone number?”	Selma K. Peterson’s phone number is 545-3899.
Has recently been asked only the question “Who has phone number 545-3899?”	[No relevant information]
[more conditions]	[information accessible relative to those conditions]

Each row of an access table consists of an elicitation condition, together with a state of mind associated with that condition. The state of mind associated with a condition reflects the information available to the subject relative to that condition.¹⁰

For example, the first row of the above table indicates that the phonebook creature is able to access the information that Selma K. Peterson’s phone number is 545-3899 in response to the question “What is Selma K. Petersen’s phone number?”. The second row indicates that creature is not able to access that information in response to the question “Who has phone number 545-3899?”. Other rows, omitted for reasons of space, are associated with yet other questions.

Note that Phone Book Creature’s access table contains many rows, even though he has a very simple method for representing his information: a single alphabetized listing and a method for searching that listing by name. So rows of an access table need not correspond to anything like modules or

¹⁰We have not specified exactly how states of mind are to be represented in the entries in the right-hand column of access tables. That is because we are trying to be neutral as between two approaches. On one approach, states of mind are represented in a way that employs fine-grained propositions. On another approach—the one we favor—states of mind are represented by sets of probability functions defined over the space of possible worlds. We think that using access tables in the second way leads to an attractive solution to the problem of logical omniscience (Hintikka 1975), a solution which is very much in the spirit of Stalnaker (1984, 1991, 1999). That is a story for another day (Elga and Rayo 2009, 2014).

subsystems in a brain. Therefore there is no reason to think that a brain must be particularly large or intricate in order to embody an access table with many rows. More generally, when we describe a subject's mental state using an access table, we specify what information is accessible to her relative to what elicitation conditions. But we do not specify the cognitive architecture underlying that pattern of access.^{11,12}

4 *Is indexing necessary?*

We have so far seen that we *can* accommodate access limitations and cognitive accomplishment by using an indexed notion of information access. But one might wonder: "Why bother? Indexing is complicated and unfamiliar. Why not just use standard off-the-shelf philosophy of mind and language tools to represent the above phenomena? Why must we introduce any kind of relativity to elicitation conditions at all?"

The answer is that (so far as we can see) the phenomena can only be accommodated by introducing relativity to elicitation conditions. The remainder of this section explains why.

Suppose that we wish to describe a subject's mental state in a way that tells us what questions the subject is able to answer. For simplicity, let us focus on fill-in-the-blank questions such as:

(α) Paris is the capital of _____.

(β) _____ is the capital of Italy.

Without using special apparatus such as indexing, how might we describe a subject's mental state in a way that settles which fill-in-the-blank questions she can answer?

¹¹For a critique of certain indexed theories which make strong claims about cognitive architecture, see Norby (2014, 34).

¹²A further clarification: access tables are meant to distinguish between cases in which someone is able to immediately deploy some information, and cases in which she isn't. But to quickly deploy some information, one needn't have explicitly represented the information in advance. So when an access table indicates that the subject has access to some information relative to an elicitation condition, this does not require that the subject has the information explicitly represented in advance.

Parikh (2008, 467) suggests a helpful analogy (Field (1978, 36–37) makes a similar point): Some book buyers care only that their books ship quickly when ordered, and not at all whether they are shipped from a pre-stocked warehouse as opposed to being printed quickly on demand. Likewise, when a theorist is interested in understanding an agent's behavior in informational terms, she might care only whether the agent is able to quickly deploy a piece of information when required, and not at all whether the information was explicitly represented in advance.

A simple approach is to describe the subject's mental state by saying what she believes. For example, when a subject believes that Paris is the capital of France, she can answer question (α). When she believes that Rome is the capital of Italy, she can answer question (β). More generally, the idea is that when a subject believes a claim, she can answer fill-in-the-blank questions corresponding to that claim (where the correspondence is as illustrated above).

That simple approach won't work. To see why, recall the case of Jack (§2.2), who can produce his neighbor's apartment number when supplied with her name, but not the other way around. Now consider the following fill-in-the-blank questions:

- (1) The person in apartment _____ is named "Beatrice Ogden".
- (2) The person in apartment 23-H is named "_____".

Because of his imperfect recall, Jack would be able to complete an exam containing only question (1) but not an exam consisting only of question (2). But the simple approach can't deliver both these results.

The simple approach faces a dilemma: Does Jack believe that the person in apartment 23-H is named "Beatrice Ogden", or not? If so, then the proposal entails that he is able to answer both questions (1) and (2). If not, then it entails that he is able to answer neither. Either way, the simple approach delivers the wrong results.

One possible fix is to say that with respect to the elicitation condition of being asked question (1), Jack believes that the person in apartment 23-H is named "Beatrice Ogden", but respect to being asked question (2), he does not. This fix amounts to a form of indexing.

An alternative fix is to introduce a *recall* relation, and claim that although Jack believes that the person in apartment 23-H is named "Beatrice Ogden", he can only recall that information in certain circumstances. On this way of thinking, belief itself is not indexed to elicitation conditions, but recall is so indexed. Note that on this way of thinking, specifying what an agent believes is not enough to settle what questions he can answer—one needs also to specify what beliefs he is able to recall relative to what conditions. So indexing figures in the overall theory of the agent's mental state as it affects his dispositions to act.

We need not object to the above proposals for present purposes. Each illustrates a way that one might employ indexing in order to characterize the mental state of an agent.

But what if one insists on avoiding indexing altogether? Can one accommodate imperfect recall just using ordinary (non-indexed, non-access-table) belief attributions, and without appealing to a special indexed recall relation?

One strategy is to introduce a very fine-grained conception of belief, and to claim that the beliefs needed to answer question (1) are different from the beliefs needed to answer question (2). One might claim, for example, that whereas correctly answering question (2) requires

believing that Beatrice Ogden lives in apartment 23-H,
correctly answering question (1) requires

believing that apartment 23-H is where Beatrice Ogden lives.

We find this proposal slightly mysterious. Our ears don't hear much of a difference between the claim that a subject believes that individual X is the F and the claim that a subject believes that the F is individual X. Furthermore, the proposal relies on the claim that Jack does not believe that Beatrice Ogden lives in apartment 23-H. But this makes it awkward to explain why he is able to easily answer the question "Does Beatrice Ogden live in 23-H?".

In reply, a friend of the above approach might stipulate that being able to correctly answer "Does Beatrice Ogden live in 23-H?" requires

believing that the answer to "Does Beatrice Ogden live in 23-H?"
is "Yes".

Notice, moreover, that this stipulation can be applied to arbitrary questions. For instance, one might say that being able to correctly answer question (2) requires

believing that the answer to the fill-in-the-blank question "The
person in apartment _____ is named 'Beatrice Ogden'" is "23-H".

and being answer to correctly question (1) requires

believing that the answer to the fill-in-the-blank question "The
person in apartment 23-H is named '_____' " is "Beatrice Ogden".

Indeed, this approach enables one to characterize the mental state of any agent, no matter what questions he is able to answer.¹³

¹³It might even allow one to characterize mental states that account for the subject's non-linguistic abilities. For example, a friend of the approach were considering could suggest that heading to the right apartment when one wishes to find Beatrice Ogden requires:

believing that in order to find Beatrice Ogden's apartment one should head to
apartment 23-H.

There are yet other alternatives. For instance, one might employ “qua” terminology, or one might repurpose modes of presentation or guises to include elicitation conditions. Or one might appeal to a context-dependent analysis of belief attributions.¹⁴

Notice that these approaches all add an extra component to each belief attribution: a component that picks out an elicitation condition. So we end up being forced to include indexing of a kind in our total theory.

5 Further applications

The above discussion suggests that an information-based theory must appeal to indexing, if it is to adequately account for a subject’s ability to answer questions. But question-answering is a kind of behavior. Therefore an information-based theory needs to appeal to indexing in order to account for behavior.

Once we adopt indexing, however, we receive additional benefits: we get an attractive way of representing fragmented, confused, and incoherent states of mind, and a pleasing representation of the difference between knowledge-that and knowledge-how. So we shall argue.

5.1 Mental fragmentation

People are sometimes in confused, incoherent, fragmented or divided states of mind. For example:

Mental Maps (cf. Lewis 1982) Whenever Andy travels to Leicester Square on foot, he treats Charing Cross Road as though it runs east-west. But whenever he travels to Leicester Square by train, he treats Charing Cross Road as though it runs north-south. And this is so even though he is confident that there is only one Queen Street near Leicester Square. *According to Andy, how is Charing Cross Road oriented?*¹⁵

Split brain (Nagel 1971) A split-brain patient is simultaneously shown a picture of a chicken in his left visual field, and a picture of a shovel in his right visual field. Upon being asked what he sees, he is disposed

¹⁴This option was suggested by Keith DeRose, but we do not know whether he endorses it in the present connection. As Jason Stanley pointed out to us, one could develop an invariantist version of the same idea.

¹⁵Evidence of fragmented mental representations in tasks relating to spatial layouts and predictions is presented in (Kalish et al. 2004, Lewandowsky and Kirsner 2000, Lewandowsky et al. 2002, Tversky 1992, 1993, 2003).

to answer: “A shovel”. But upon being given a toy chicken and a toy shovel and asked to hand over an object similar to what he sees, he is disposed to hand over the chicken. *According to the patient, what picture was he shown?*

Disavowed Racism (Schwitzgebel 2010, 532) Juliet “has critically examined the literature on racial differences in intelligence, and she finds the case for racial equality compelling. She is prepared to argue coherently, sincerely, and vehemently for equality of intelligence and has argued the point repeatedly in the past... And yet Juliet is systematically racist in most of her spontaneous reactions, her unguarded behavior, and her judgments about particular cases. When Juliet is on the hiring committee for a new office manager, it won’t seem to her that the black applicants are the most intellectually capable, even if they are.” *According to Juliet, how do the average intelligence levels of blacks and whites compare?*¹⁶

Witness When asked to pick the perpetrator out of various lineups, a witness is disposed to confidently pick men with a particular sort of thin, bearded face. But when asked to produce a sketch of the perpetrator with the help of a police sketch artist, that same witness is disposed to confidently produce a sketch of a fat, beardless face. *According to the witness, what did the perpetrator look like?*

Each of these people is in a divided state of mind—none of them has a completely coherent conception of how the world is. How should we represent the mental states of such people? We would like to have the mental state we ascribe to an agent lead to appropriate predictions about how the agent is disposed to act. But the simplest such proposals fail to do so.¹⁷

For example, consider the case of the witness. Suppose that the witness is sitting at home, not thinking about the crime at all. How should we represent her mental state in order to capture the full range of her dispositions to action (both her pick-out-of-a-lineup dispositions and her create-a-sketch dispositions)?

If one simply said “The witness **believes** that the perpetrator had a beard”, that would lead to the incorrect prediction that the witness is disposed to produce sketches of bearded men. If one simply said “The witness

¹⁶Schwitzgebel (2010) calls the attitude exhibited in this and similar cases “in-between belief” and convincingly refutes attempts to analyze it in terms of existing categories such as pure belief or disbelief.

¹⁷The remainder of this subsection draws heavily on Schwitzgebel (2010).

disbelieves that the perpetrator had a beard”, that would lead to the incorrect prediction that the witness is disposed to pick beardless men out of lineups (Schwitzgebel 2010, 543).

How about saying “The witness is highly **uncertain** whether the perpetrator has a beard,” or “The witness **suspends judgment** on the matter”? This would also lead to incorrect predictions. It would predict that the witness is disposed to vacillate, exhibit doubts, or refuse to deliver a strong judgment when queried about the appearance of the perpetrator. But in fact she is disposed to confidently pick bearded men from lineups and to confidently produce sketches of beardless men.

Now suppose one simply said “The witness is **confused** about whether the perpetrator has a beard,” or “The witness **both** believes that the perpetrator is bearded and believes that the perpetrator is beardless.” The resulting proposal would be at best incomplete. To see why, consider a second witness. The second witness is just like the first one, except that he is disposed to pick *beardless* men out of lineups, but to produce sketches of *bearded* ones. In other words, the second witness has a pattern of dispositions that is “swapped” with respect to the first.

The two witnesses have different dispositions concerning the beardedness of the perpetrator. But the description “confused about whether the perpetrator has a beard” applies to the two witnesses equally well. And the same goes for the description “believes that the perpetrator is bearded and believes that the perpetrator is beardless”. So neither description captures the difference in dispositions between the two witnesses.

5.2 *How to characterize fragmented mental states*

We have been considering the question: How should we represent the witness’s mental state in order to sufficiently capture her dispositions to action? We have seen that various simple proposals do not succeed. But a proposal relying on indexed mental states is well suited for the job, as is illustrated by the following access table:

<i>Elicitation Condition</i>	<i>Information accessible relative to that condition</i>
Visual recognition of perpetrator	Perpetrator had a thin, bearded face
Production of sketches of perpetrator	Perpetrator had a fat, un-bearded face

This access table gets across that the witness is confused (i.e., is incoherent, fragmented, of two minds) about whether the perpetrator had a beard. More importantly, it pins down the specific type of confusion. It specifies circumstances in which the witness is disposed to act as if the perpetrator was bearded, and circumstances in which she is disposed to act as if the perpetrator was beardless. That specification of circumstances matters, because it increases the extent to which the table captures the witness's behavior.¹⁸

5.3 Indexing and abilities

We have seen how to use indexing to represent the mental states of agents who have imperfect memory access and agents who confused or incoherent. Indexing can also represent the mental states of agents whose propositional knowledge differs from their know-how. Let us explain.

Sometimes, in order to make use of some information, a subject must perform the right sort of *action*. For example:

Slot A patient suffering from visual agnosia is presented with a surface with a slot cut into it. The slot is oriented diagonally. When given a card and asked to orient the card at the same angle as the slot, or to verbally report the orientation of the slot, the patient is unable to comply. But when asked to *place* the card in the slot, she immediately orients the card correctly as he extends her hand to the slot.¹⁹

¹⁸Gendler (2008a,b) has brought attention to some fascinating cases involving divided states of mind. These cases, it is argued, highlight the difference between belief and a state of mind dubbed "Alief". For example, Gendler (2008b, 636) describes a traveller who leaves home without her wallet and asks a friend for a loan. When the traveller receives the money and needs a place to store it, she unreflectively reaches for her wallet. Such cases exhibit a distinctive asymmetry. One behavior—asking a friend for money—has features which are paradigmatically rational (they are responsive to evidence, for example). But another behavior—searching for one's wallet in an effort to store the money—does not.

Access tables can be used to represent some aspects of cases of this kind. In particular, there is a simple access table that captures the wallet-related dispositions to behavior in the example above. On the other hand, some cases of Alief involve dispositions that are not particularly robust, and some of them involve behaviors that are not naturally rationalizable. Access tables may not be well-suited to representing such cases.

In using an access table to represent a divided state of mind, one doesn't settle whether that case exhibits the distinctive asymmetry mentioned above. Nor does one take sides on whether the cases discussed in Gendler (2008a,b) form a psychological natural kind.

¹⁹Here is a description of one such patient, known in the psychology literature as "D.F." (Milner et al. 1991, 418–421; Milner 1997, 1250):

Even though [D.F.] was very poor at describing or demonstrating the orientation of a line or slot, she could still reach out and post a card into the same slot

In this case, the orientation of the slot is available to the patient for performing one sort of action (putting the card in the slot) but not another sort (just orienting the card).

As before, we can represent the patient's limited access to her information with an appropriate table. In this case, the elicitation conditions differ not in what question the subject is asked, but rather in the way the subject indicates her answer:

<i>Elicitation Condition</i>	<i>Information accessible relative to that condition</i>
Indicate orientation of slot by orienting card	[Uncertain about orientation of slot]
Indicate orientation of slot by inserting card	Slot is diagonal

The above example involves a visual deficit, but similar cases are completely mundane. For example, can you immediately say whether the following claim is true?—

Your standard shoe-tying method involves looping the cord around your right thumb.

Some of us have trouble immediately answering this question. Even so, the requisite information is in us somewhere—otherwise we wouldn't be able to tie our shoes. What causes the trouble is that the information is encoded as part of a motor memory sequence, and an actual or imagined shoe-tying operation is necessary to extract it.

In the shoe-tying example, the subject has access to some information for the purpose of performing one sort of action but not another. It can also happen that a subject has access to some information for the purpose of performing one sort of action, and access to *contrary* information for the purpose of performing another sort.

For example, consider the following surprising fact about cycling: in order to initiate a slow right turn one typically leans to the *left*.²⁰ All competent cyclists are in a position to deploy that information for the purposes of

without error. Similarly, despite being unable to report (verbally or manually) the width of a rectangular block, she would still tailor her finger-thumb grip size perfectly in advance of picking it up. In short, she could guide her movements using visual cues of which she seemed completely unaware. (Goodale and Milner 2006, 660)

We learned of the existence of such patients from Kelly (2000, 171).

²⁰See <http://socrates.berkeley.edu/~fajans/Teaching/Steering.htm>. This fact was brought to our attention by Richard Holton.

making turns. But few are able to deploy it for the purposes of explaining how to ride a bike. In fact, when giving verbal explanations, most cyclists are disposed to report exactly the opposite.²¹ In this case, too, we can represent the cyclist's differential access to information with an appropriate access table.

5.4 *Knowledge-how versus knowledge-that*

Natural language sometimes distinguishes between knowledge-how and knowledge-that. Ordinary ascriptions of knowledge-how tend to track the subject's ability to perform practical tasks competently. In contrast, ascriptions of knowledge-that tend to track the subject's ability to articulate information linguistically, or use it in reasoning. Knowing *how* to ride a bicycle, for example, ordinarily goes with having the ability to competently ride a bicycle. Knowing *that* the Battle of Hastings occurred in 1066, on the other hand, ordinarily goes with the ability to state when the Battle of Hastings took place, or use that information in the course of figuring out whether the Battle of Hastings preceded the Battle of Waterloo.

Appealing to indexed mental states does not on its own deliver a characterization of the knowledge-that/knowledge-how distinction. But it improves our understanding of that distinction by locating it in a broader context. For notice how natural it is to describe the above examples in terms of knowledge-how and knowledge-that:

- The ordinary shoe-tying practitioner knows how to tie his shoes. But he doesn't know that his shoe-tying method involves looping around his right thumb.
- The bicycle rider knows how to perform a slow right turn. But he doesn't know that to perform a slow right turn one must start by leaning left.

In each case, saying that the subject has knowledge-how conveys that the subject has access to relevant information for the purpose of performing one sort of action (tying shoelaces, *putting the card in the slot*, riding a bicycle). In contrast, saying that the subject lacks knowledge-that conveys that the subject lacks access to that information for the purpose of performing another sort of action (explaining shoelace-tying, explaining bicycle riding).

This suggests that in general, the difference between having knowledge-that and having knowledge-how amounts to the difference between having

²¹Cf. Stalnaker (1999, 263).

information available for one sort of action, and having it available for another. Indexing is well suited to representing this sort of difference.^{22,23}

6 Frequently asked questions

Why don't you give a more detailed theory, and specify the actual cognitive mechanisms which are responsible for limitations of information access?

We would love to have a detailed computational theory capable of explaining all of the cases discussed. But we don't have one, and we think that the search for such a theory is still in its infancy.²⁴

Since access tables don't settle what a subject's mental architecture is, what is the point of introducing them?

Even though we do not have a full theory of the mechanisms leading to limitations in human information access, it is still helpful to have a unified, systematic framework for displaying those limitations. Such a framework helps us sharpen up and draw attention to the phenomenon.

Notice also that giving a subject's access table can help us understand her behavior even in the absence of a complete theory of mental representation. To see how, recall the cases from the first paragraph of this paper: the soldier who evades mines on a mine-field, the student who writes down correct answers, and the bank manager who opens a combination safe. In explaining these successes, it is extremely helpful to speak in terms of information possession. When we say that the soldier succeeded in evading the mines because he had access to information about their location, we shed light on his behavior even if we don't know exactly how his brain works. For

²²There is a much-disputed question: can knowledge-how be analyzed in terms of knowledge-that? (Glick 2011, Hawthorne and Stanley 2008, Ryle 1949, Stanley 2011, Stanley and Williamson 2001). Even those who think that knowledge-how can be analyzed in terms of knowledge-that agree that typical ascriptions of knowledge-how make systematically *different* claims than typical ascriptions of knowledge-that. So it is open to them to represent that difference in terms of information being available for different purposes. For example, according to Stanley and Williamson (2001) knowing how to ride a bicycle amounts to knowing, of a particular way *w* thought of under a "practical" mode of presentation, that *w* is a way to ride a bicycle. It is compatible with this analysis to say that having that information under a practical mode of presentation amounts to having it available for the purpose of performing certain actions.

²³For more on using indexing to represent know-how, see Elga (2012), Rayo (2013, §4.5), and Bianchi (2014).

²⁴Promising early work in this direction includes Egan (2008).

comparison, think about how awkward it would be to explain the soldier's success purely in terms of fundamental physical laws.

Similarly, it is useful to speak in terms of purpose-relative access to information even when we're not sure exactly how information is represented.

*Even granting that talking in terms of information in general has some explanatory value, your way of doing so seems particularly empty. An access table is just a big list of dispositions to behave. So by appealing to sufficiently fine-grained elicitation conditions, we can cook up an access table to match any pattern of behavior. How could such an exercise be anything more than trivial?*²⁵

We admit that access tables would not be interesting if in order to use them, one needed to resort to the trick of appealing to arbitrarily specific elicitation conditions. But one needn't resort to that trick. For example, the access table for the "dreamt" case contained comparatively coarse-grained and natural conditions such as "The word 'dreamt' has been recently made salient". It did not contain unnatural and exceedingly fine-grained conditions such as "The word 'dreamt' was made salient on a Tuesday, with a voice of frequency 129.2389 Hz, when the subject was hungry, . . .".

Does this paper provide an analysis of belief, or a semantics for belief attributions?

No. The present focus is rather in advocating a way of describing mental states that can accommodate cases of limited recall and deductive abilities. The resulting notion of an access table does not exactly match up with the ordinary folk notion of belief, and it would take work to connect the two notions.²⁶

Furthermore, notice that ordinary belief attributions do not on their own allow us to easily accommodate all cases of limited information access. For example, recall the case of Jack, who is able to fill in the blank in this question:

(1) The person in apartment _____ is named "Beatrice Ogden".

but not this one:

(2) The person in apartment 23-H is named "_____".

Does Jack believe that the person in apartment 23-H is named "Beatrice Ogden"? However we answer, it would take extra work to explain how that answer would predict the ability to answer (1) but not (2).

²⁵Thanks here to Andy Egan.

²⁶Norby (2014, 34–35) raises some challenges for those who wish to use an indexed-based framework to analyze belief.

In sum: we don't give an analysis of belief. But doing so is neither necessary nor sufficient for shedding light on cases involving limited recall and mental fragmentation.

References

Dylan Bianchi. Know how: Information you can use. Manuscript, August 2014.

David Braddon-Mitchell. *The Philosophy of Mind and Cognition*. Blackwell Pub., 2007.

Bryson Brown and Graham Priest. Chunk and permeate, a paraconsistent inference strategy. part i: The infinitesimal calculus. *Journal of Philosophical Logic*, 33(4):379–388, 2004.

Yuri Cath. The ability hypothesis and the new knowledge-how. *Noûs*, 43(1): 137–156, 2009.

Mark Crimmins. Tacitness and virtual beliefs. *Mind and language*, 7(3): 240–263, 1992.

Andy Egan. Seeing and believing: Perception, belief formation and the divided mind. *Philosophical Studies*, 140(1):47–63, 2008.

Adam Elga. Belief fragmentation. 2012. Yale University (Graduate Seminar Visit).

Adam Elga and Agustín Rayo. Fragmented belief states and logical omniscience. 2009. CUNY Graduate Center presentation.

Adam Elga and Agustín Rayo. Fragmentation and logical omniscience. 2014. Manuscript.

Hartry Field. Mental representation. *Erkenntnis*, 13(July):9–18, 1978.

Tamar Szabó Gendler. Alief in action (and reaction). *Mind and Language*, 23 (5):552–585, 2008a.

Tamar Szabó Gendler. Alief and belief. *Journal of Philosophy*, 105(10):634–663, 2008b.

Ephraim Glick. Two methodologies for evaluating intellectualism. *Philosophy and Phenomenological Research*, 83(2):398–434, 2011.

- Mel Goodale and David Milner. One brain—two visual systems. *The Psychologist*, 19(11):pp. 660–663, 2006.
- John Hawthorne and Jason Stanley. Knowledge and action. *Journal of Philosophy*, 105(10):571–590, 2008.
- Jaakko Hintikka. Impossible possible worlds vindicated. *Journal of Philosophical Logic*, 4(4):475–484, 1975.
- Michael L. Kalish, Stephan Lewandowsky, and John K. Kruschke. Population of linear experts: Knowledge partitioning and function learning. *Psychological Review*, 111(4):1072–1099, 2004.
- Sean D. Kelly. Grasping at straws: Motor intentionality and the cognitive science of skillful action. In *Essays in Honor of Hubert Dreyfus, Vol. II*. MIT Press, 2000.
- Stephan Lewandowsky and Kim Kirsner. Knowledge partitioning: Context-dependent use of expertise. *Memory & Cognition*, 28(2):295–305, 2000.
- Stephan Lewandowsky, Michael Kalish, and S.K. Ngang. Simplified learning in complex situations: Knowledge partitioning in function learning. *Journal of Experimental Psychology*, 131(2):163–193, 2002.
- David Lewis. Logic for equivocators. *Nous*, 16:431–441, 1982.
- A.D. Milner. Vision without knowledge. *Phil. Trans. R. Soc. Lond. B*, 352: 1249–1256, 1997.
- A.D. Milner, D.I Perrett, R.S. Johnston, P.J. Benson, T.R. Jordan, D.W. Heeley, D. Bettucci, F. Mortara, R. Mutani, E. Terazzi, and D.L.W. Davidson. Perception and action in ‘visual form agnosia’. *Brain*, 114(1):405–428, 1991.
- Thomas Nagel. Brain bisection and the unity of consciousness. *Synthese*, 22: 396–413, 1971.
- Aaron Norby. Against fragmentation. *Thought*, 2(4):30–38, 2014.
- Rohit Parikh. Sentences, belief and logical omniscience, or what does deduction tell us? *Review of Symbolic Logic*, 1(4):459–476, 2008.
- Lawrence Powers. Knowledge by deduction. *Philosophical Review*, 87(3): 337–371, 1978.

- Augustin Rayo. *The Construction of Logical Space*. Oxford University Press, 2013.
- G. Ryle. *The Concept of Mind*. University of Chicago Press, Chicago, IL, 1949.
- Eric Schwitzgebel. Acting contrary to our professed beliefs, or the gulf between occurrent judgment and dispositional belief. *Pacific Philosophical Quarterly*, 91:531–553, 2010.
- Robert Stalnaker. *Inquiry*. MIT Press, Cambridge, Mass., 1984.
- Robert Stalnaker. The problem of logical omniscience. *Synthese*, 89(3):425–440, 1991.
- Robert C. Stalnaker. The problem of logical omniscience, II. In Robert C. Stalnaker, editor, *Context and Content*. Oxford University Press, Oxford, 1999.
- Jason Stanley. Knowing (how). *Noûs*, 45(2):207–238, 2011.
- Jason Stanley and Timothy Williamson. Knowing how. *Journal of Philosophy*, 98(8):411–444, 2001.
- Barbara Tversky. Distortions in cognitive maps. *Geoforum*, 23(2):131–138, 1992.
- Barbara Tversky. Cognitive maps, cognitive collages, and spatial mental models. In A.U. Frank and I. Campari, editors, *Spatial Information Theory: A Theoretical Basis for GIS, Proceedings COSIT '93*. Springer, Berlin, 1993.
- Barbara Tversky. Structures of mental spaces: How people think about space. *Environment and Behavior*, 35(1):66–80, 2003.