

Fragmentation and information access*

Adam Elga and Agustín Rayo[†]

October 29, 2019

Abstract

Specifying the information possessed by a subject is not enough to account for her behavior. One must also specify the conditions under which information is *accessible* to her. Attention to access conditions allows us to model agents who exhibit imperfect recall, confusion, and mental fragmentation. It also sheds light on the difference between propositional knowledge and knowledge-how.

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. We can account for each of these successes by noting 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 ascriptions of information to account for the behaviors manifested in such cases, they must be “indexed”: they must determine what information is available to the agent *relative to various circumstances*.²

*Revised version to appear in *The Fragmented Mind*, Cristina Borgoni, Dirk Kindermann, and Andrea Onofri, eds. Oxford University Press. The authors of this chapter contributed equally to it. Thanks to Andrew Bacon, Ross Cameron, David Chalmers, Jonathan Cohen, Keith DeRose, Cian Dorr, Hartry Field, Branden Fitelson, Jeremy Goodman, Eric Hubble, David Hunter, Frank Jackson, Shivaram Lingamneni, John MacFarlane, Michael Rescorla, Ted Sider, 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

¹Cf. Braddon-Mitchell and Jackson (2007).

²The suggestion in Stalnaker (1984) that logical omniscience failures can be understood

2 *Limitations on information access: 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 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.³

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 (follow-

in terms of fragmented belief states was the core motivation for the present model. Braddon-Mitchell and Jackson (2007, 199–200) also uses fragmented coarse-grained belief states to accommodate failures of logical omniscience. Yalcin (2008, Ch. 3), Yalcin (2015), and Yalcin (2016) develop that same suggestion, proposing an elegant model on which all-or-nothing belief is relative to questions, understood as partitions of logical space. The treatment of logical omniscience failures in those works uses privileged partitions to represent which propositions are accessible to an agent, and so differs from the present treatment. (See especially Yalcin (2016, n. 26).) Egan (2008a) endorses a treatment of fragmented credences similar to the present one and interestingly suggests that mental fragmentation might be practically indispensable for agents with perceptual belief forming mechanisms anything like human ones—mechanisms that are less than perfectly reliable but which nevertheless produce immediate belief in certain circumstances. Other work congenial to indexing includes Brown and Priest (2004), Greco (2015, 2019), Koralus (2014), Lewis (1982), Parikh (2008), Schwitzgebel (2001), and, of course, many of the articles in the present volume.

³Talk 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.

ing Stalnaker (1991)) 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.⁴

3 *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 Jack is instead asked "How do you know Beatrice Ogden?", he 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.

4 *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"?

Think about this puzzle for a moment before reading on.

⁴Indeed, as the designers of computer databases are well aware, this is not the exception but the rule. Setting up a database to make its information easily accessible for one purpose will often make that information less accessible for other purposes.

Now for the metapuzzle. Even before hearing the puzzle you knew perfectly well that “dreamt” is a word of English and knew how to spell it. Indeed, you knew that “dreamt” is a word of English ending in “MT”. (If before hearing the puzzle someone had excitedly announced to you, “‘Dreamt’ is an English word ending in ‘MT’!” you might have thought to yourself, “I knew *that* already.”)⁵

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 ‘MT’?”^{6,7}

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.

5 Access tables

The above example suggests that in order to account for the behavioral dispositions of a subject with imperfect access to her information, we need to specify what information is available to an agent relative to various purposes. Here is a simple way of doing so.

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 given a name, and asked for the apartment number where the named person lives. Another one is being given an apartment number and asked for the name of the person who inhabits the named apartment. Elicitation conditions often prompt an agent to deploy

⁵The “I knew that already” argument, and the puzzle that it concerns, is based on the “deny” example and discussion in Powers (1978, 341). See also the “factoring” example from Stalnaker (1991, 438), Crimmins (1992), and the “iceman” example from Egan (2008b, 51).

⁶Philosophy 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).

⁷Some 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 (Elga and Rayo 2009, 2019a, Stalnaker 1984, 1991), but won’t argue the point here.

information for a particular purpose, so it is sometimes convenient to think of elicitation conditions as purposes.

Below is an “access table” for Jack—a table specifying what information is available relative to each elicitation condition. 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>
Asked for an apartment number given a name	The person in “23-H” is named “Beatrice Ogden”
Asked for a name given an apartment number	[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.⁸

For example, the first row of the above table entails that Jack is able to access the information that the person in the apartment labeled “23-H” is named “Beatrice Ogden” in directing a letter addressed “To: Beatrice Ogden”. The second row entails that Jack is not able to access that information in response to the question “What is the name of the person in 23-H?”.

Three clarifications about access tables are in order.

First, a subject’s access table leads to predictions about her rational actions, given suitable information about her desires. Roughly speaking, the subject should act so as to satisfy her desires, given the information she has access to relative to her present elicitation condition.⁹

⁸We 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 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 (Elga and Rayo 2019b). We think (Elga and Rayo 2009, 2019a) 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) and Lewis (1982).

⁹The above link is appropriate under the idealization of all-or-nothing information possession. In fact we think that it is better to work in terms of credences rather than

Second, access tables are representations of an agent's informational state. But when we focus on a special class of agents—agents who use their information to further their goals, in the right sort of way—access tables can also represent an agent's behavioral dispositions. Such agents behave as their access tables predict. We can therefore say that an access table “accounts for” the behavioral dispositions of such an agent just in case it predicts them.

Third, access tables are meant to distinguish between cases in which someone is able to immediately deploy some information, and cases in which she is 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) offers a helpful analogy: 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.¹⁰

6 *Individuating elicitation conditions*

One way for a theory to be explanatory is for it to identify patterns and show that relevant facts are instances of those patterns. Access tables are explanatory in at least this sense. By specifying the subject's information relative to various purposes, they supply a systematic description of her rational dispositions.

This assumes, however, that the access table's elicitation conditions are not individuated too finely. Otherwise, an access table might become a mere listing of overly specific dispositions, and so fail to provide useful explanations of behavior.¹¹ So we are left with a pressing question: how finely should elicitation conditions be individuated?

all-or-nothing information possession, and so to take an agent's decision-theoretic state to be an access table whose rows each contain a credence function. If we do that, the link between access tables and rational action will be: for an agent in choice condition c , the rational actions are exactly that maximize expected utility relative to the credence function on row c of her access table (Elga and Rayo 2019b).

¹⁰Field (1978, 36–37) makes a similar point.

¹¹Norby (2014, §2), Hoek (2019, 137–139), Quilty-Dunn and Mandelbaum (2018, 2358–2359).

One might answer this question by deferring to cognitive science: by requiring that the structure of an access table match the structure of the subject's cognitive architecture. We do not appeal to cognitive science in this way. Instead, we say that access tables are neutral about the details of the cognitive architecture that is responsible for the subject's patterns of information access, just as orthodox decision theory remains neutral about the details of the cognitive architecture that is responsible for the agent's degrees of belief.¹² This means, in particular, that the rows of an access table aren't intended to map cleanly onto components of a realistic cognitive psychology. For example, there is no suggestion that different rows correspond to different cognitive modules.¹³

To repeat: attributing an access table to an agent does not settle what is it about the agent's mental architecture that links certain pieces of information to certain conditions. Access tables are intended to settle what information an agent possesses, but not the mechanism through which that possession is accomplished. This facilitates one valuable role for information: as a "common currency" that might be shared even between creatures whose architectures differ radically. For example, consider a cantina frequented by drinkers of highly diverse alien species (Wookieepedia 2019). We may wish to say that two very different drinkers possess the same bit of information relative to a particular condition. Making an appropriate claim about the drinkers' access tables would accomplish this even if the drinkers have radically different cognitive architectures.¹⁴

So how should one determine the fineness of grain of an access table's elicitation conditions? The answer is that the set of choice situations should strike a good balance between two competing requirements. First requirement: elicitation conditions are fine-grained enough to account for as many of the subject's rational dispositions as possible. Second requirement: elicitation conditions are coarse-grained enough to ensure that the resulting access

¹²Here we follow the common, though not universal, practice of presupposing a "dispositionalist" reading of credence functions, as opposed to a reading on which the attribution of a credence function makes a detailed claim about the cognitive architecture that realizes it. For example, we assume that attributing an agent credence 1/3 in the claim that it rains tomorrow does not entail that the agent has an explicit mental representation with the structure of the sentence "it will rain tomorrow" tagged with the number "1/3". For arguments against the more cognitively committal reading of credence attributions, see Byrne (2019), Holton (2014).

¹³For a worthy defense of mental fragmentation from the perspective of philosophically-informed cognitive science and psychology, see Mandelbaum and Bendana (Forthcoming).

¹⁴We don't mean to suggest that questions about cognitive architectures—in particular, the human cognitive architecture—are unimportant. We consider them tremendously important. They are just outside the scope of the present project.

table delivers a reasonably systematic description of the subject's rational dispositions.¹⁵

7 *Is indexing necessary?*

We have seen that “indexing”—i.e. relativity to elicitation conditions—can be used to account for the behavioral dispositions of a subject with imperfect access to her information. But one might wonder: “Is indexing really necessary? Why not just use off-the-shelf tools from the philosophy of mind to account for such dispositions? Why must we introduce any kind of relativity to elicitation conditions at all?” The answer is that (so far as we can see) indexing is indispensable. 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 a special apparatus similar to indexing, how might we describe a subject's mental state in a way that settles which fill-in-the-blank questions she can answer?

A simple approach is to describe the subject's mental state by saying what she believes, using the ordinary notion of belief. 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 in general. To see why, go back to the case of Jack (§3), who can produce his neighbor's apartment number when

¹⁵There are several strategies for ensuring a good balance between these requirements. On a realist approach, one insists that there is a single objective fact of the matter about what access table represents an agent's mental state. One might then attempt to characterize the relevant access table by a process of radical interpretation (perhaps by adapting Lewis (1974) to include the two requirements above). On an alternative “modeling” approach, one makes no claims about what access table a given agent “really” has. Instead one treats the ascription of an access table as a theoretical move to be judged by its theoretical fruitfulness. On that way of thinking, the above two requirements are thought of not as conditions that the one true theory of access-table-determination satisfies, but rather as guides to theorists for producing fruitful models of fragmented agents.

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 containing only question (2). But the simple approach can't straightforwardly deliver both these results.¹⁶

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.

Can the simple approach be fixed?

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 with respect to being asked question (2), he does not. This fix amounts to a form of indexing.¹⁷

Another possible fix is to take beliefs to be highly unstable. For example, one could claim that before Jack is asked any questions, he does not believe

¹⁶Compare to Schwitzgebel (2001, 79).

¹⁷There are other ways of implementing the same general idea. For example, one might use the notion of a proposition-under-a-mode-of-presentation. Modes of presentation are a form of indexing, since a proposition-under-a-mode-of-presentation can be thought of as a pair of a proposition and a condition with respect to which the proposition might be accessed. In this case the index attaches to the object of belief, rather than the belief relation itself, but we do not take this difference to be theoretically significant. As far as we can tell, propositions-under-a-mode-of-presentation can do the same work as access tables, when suitably spelled out. (For further discussion, see Rayo 2013, §4.2.)

Proposals based on guises, or "qua" terminology also depend on a kind of indexing. An alternative approach is to claim 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".

but 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".

This approach adds an extra component to each belief attribution: a component that picks out an elicitation condition. So it too introduces a form of indexing. (We are grateful to Keith DeRose for bringing up related ideas in discussion. As Jason Stanley pointed out to us, invariantist versions of the proposals are also available.)

that the person in apartment 23-H is named “Beatrice Ogden”. But he would quickly acquire that belief under certain circumstances (such as being asked question (1)).¹⁸

To an advocate of this instability fix we pose a question: what is it about Jack’s pre-question mental state that explains why, upon being asked question (1), he so quickly and effortlessly comes to realize the answer? The advocate might answer: “There is nothing about Jack’s pre-question mental state that explains this”. In that case we reply that the resulting notion of mental state is impoverished, since it delivers the unwelcome result that a subject’s memories need not be reflected in her mental state.

Alternatively, the advocate of the instability fix might take memories to be part of the subject’s mental state, but propose that an agent’s stock of memories is distinct from her beliefs: a memory only becomes a belief when it is recalled. For example, Jack starts out with the memory that the person in apartment 23-H is named “Beatrice Ogden”, but not the corresponding belief. That memory only becomes a belief when Jack hears question (1).

This proposal is form of indexing. Notice, in particular, that Jack’s mental state can only settle what questions he is able to answer by settling what memories he is able to recall *relative to what conditions*. So indexing figures in the overall theory of Jack’s mental state.

Moral: There are many different ways in which indexing might be used to account for the behavioral dispositions of an agent with imperfect access to her information. But it is hard to see how indexing can be avoided altogether.

8 Access tables and belief

How do an agent’s beliefs relate to an indexed representation of her mental state? Here is a natural suggestion:

ROW BELIEF For an agent to believe that p is for her to have access to the information that p relative to some elicitation condition.

This suggestion is tempting, but *we do not endorse it*.¹⁹

¹⁸Compare to Audi (1994) and to the “shifting view” endorsed in Rowbottom (2007) and criticized in Schwitzgebel (2010, 543). A version of this proposal was independently brought to our attention by Bacon (2017).

¹⁹Lewis (1996, 565) tentatively endorses an analogous principle for knowledge: “S knows that P iff any one of S’s compartments knows that P .” Stalnaker (1984, 83) relies on the premise that “what it means to say that an agent believes that P at a certain time is that some one of the belief states the agent is in at that time entails that P ” to show that “[i]t is compatible with the pragmatic account that the rational dispositions that a person has at one time should arise from several different belief states”.

What's the alternative? Although we think that the relationship between an agent's beliefs and her access table is too complex and messy to be captured by any simple principle, the following is a useful heuristic:

DISPOSITIONAL BELIEF An ordinary belief ascription is associated with a cluster of dispositions: the dispositions to act and react "in patterns that ordinary people would regard as characteristic of having that [belief]." (Schwitzgebel 2001, 75)

It is appropriate to ascribe a belief to a fragmented subject if and only if the subject's access table predicts a sufficient portion of the dispositions associated with that belief.

We develop this idea further in Elga and Rayo (2019a).

9 *Mental fragmentation*

We have seen that access tables can be used to account for the behavioral dispositions of agents with limited access to their information. We will now argue that they can also be used to account for the dispositions of agents in in confused, incoherent, fragmented or divided states of mind.

Consider the following examples:

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?*²⁰

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 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?*

²⁰Evidence 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).

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?*²¹

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.²²

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 account for 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 **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

²¹Schwitzgebel (2010) calls the attitude exhibited in this and similar cases “in-between belief” and convincingly argues against attempts to analyze it in terms of existing categories such as pure belief or disbelief.

²²The remainder of this subsection draws heavily on Schwitzgebel (2010).

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.

One might say “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.

In contrast, that the witness’s dispositions can be account for using 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 can account for the witness’s

behavior.^{23,24}

10 *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.²⁵

²³Gendler (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 accounts for 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.

²⁴Giving a fragmented treatment of the above cases also removes pressure to say that such cases are violations of so-called “KK” and “BB” principles in epistemology, a result that some authors find desirable (Greco 2015).

²⁵Here 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 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

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	[No information 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 counterintuitive fact about cycling: to turn right, you typically start by steering left.²⁶ All competent cyclists are in a position to deploy that information for the purposes of making turns. But few are able to deploy it for the purposes of explaining how to ride a

Milner 2006, 660)

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

²⁶See Fajans (2000, 2015). We first learned about counter-steering from Richard Holton, though the phenomenon is well-known in motorcycling circles.

bike. In fact, when giving verbal explanations, most cyclists are disposed to report exactly the opposite.²⁷ In this case, too, we can represent the cyclist's differential access to information with an appropriate access table.²⁸

References

- Robert Audi. Dispositional Beliefs and Dispositions to Believe. *Noûs*, 28(4): 419, December 1994. ISSN 00294624. doi: 10.2307/2215473.
- Andrew Bacon. Response. Presented at Pacific APA session on Elga and Rayo's "Fragmented Decision Theory", April 2017.
- Dylan Bianchi. Know how: Information you can use. Manuscript, August 2014.
- David Braddon-Mitchell and Frank Jackson. *The Philosophy of Mind and Cognition, Second Edition*. Blackwell, 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.
- Alex Byrne. Perception and probability, May 2019.
- Yuri Cath. The ability hypothesis and the new knowledge-how. *Noûs*, 43(1): 137–156, 2009.

²⁷Cf. Stalnaker (1999, 263).

²⁸There 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). The preceding discussion suggests that the difference between having knowledge-that and having knowledge-how might amount to the difference between having information available for one sort of action, and having it available for another. Indexing is well suited to representing this sort of difference.

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.

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

- 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, 2008a.
- Andy Egan. Seeing and believing: perception, belief formation and the divided mind. *Philosophical Studies*, 140(1):47–63, July 2008b. ISSN 0031-8116, 1573-0883. doi: 10.1007/s11098-008-9225-1.
- 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. 2019a. Manuscript.
- Adam Elga and Agustín Rayo. Fragmented decision theory. 2019b.
- J. Fajans. Steering in bicycles and motorcycles. *American Journal of Physics*, 68 (7):654–659, July 2000. ISSN 0002-9505, 1943-2909. doi: 10.1119/1.19504.
- J. Fajans. The physics of bikes, April 2015. URL <https://www.youtube.com/watch?v=OPsFU9e3AD0&feature=youtu.be>.
- 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.
- Daniel Greco. Iteration and fragmentation. *Philosophy and Phenomenological Research*, 91(3):656–673, 2015. ISSN 1933-1592. doi: 10.1111/phpr.12086. URL <http://dx.doi.org/10.1111/phpr.12086>.
- Daniel Greco. Fragmentation and higher-order evidence. In Mattias Skipper and Asbjørn Steglich-Petersen, editors, *Higher-Order Evidence: New Essays*. Oxford University Press, Oxford, New York, October 2019. ISBN 978-0-19-882977-5.

- 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.
- Daniel Hoek. *The Web of Questions*. PhD thesis, New York University, 2019.
- Richard Holton. Intention as a model for belief. In Manuel Vargas and Gideon Yaffe, editors, *Rational and Social Agency: Essays on the Philosophy of Michael Bratman*. Oxford University Press, 2014.
- 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.
- Philipp Koralus. The erotetic theory of attention: Questions, focus and distraction. *Mind & Language*, 29(1):26–50, February 2014. ISSN 02681064. doi: 10.1111/mila.12040.
- 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. Radical interpretation. *Synthese*, 23:331–344, 1974.
- David Lewis. Logic for equivocators. *Nous*, 16:431–441, 1982.
- David K. Lewis. Elusive knowledge. *Australasian Journal of Philosophy*, 74(4): 549–567, 1996. doi: 10.1080/00048409612347521.
- Eric Mandelbaum and Bendana. The fragmentation of belief. In *The Fragmentation of Mind*. Oxford University Press, Forthcoming.
- 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.
- Jake Quilty-Dunn and Eric Mandelbaum. Against dispositionalism: belief in cognitive science. *Philosophical Studies*, 175(9):2353–2372, September 2018. ISSN 1573-0883. doi: 10.1007/s11098-017-0962-x.
- Agustín Rayo. *The Construction of Logical Space*. Oxford University Press, June 2013. ISBN 978-0-19-966262-3. doi: 10.1093/acprof:oso/9780199662623.001.0001.
- Darrell P. Rowbottom. 'In Between Believing' and Degrees of Belief. *Teorema: International Journal of Philosophy*, 26(1):131–137, 2007.
- G. Ryle. *The Concept of Mind*. University of Chicago Press, Chicago, IL, 1949.
- Eric Schwitzgebel. In-between Believing. *The Philosophical Quarterly*, 51(202): 76–82, 2001. ISSN 1467-9213. doi: 10.1111/1467-9213.00215.
- 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.

Wookieepedia. Cantina. <https://starwars.fandom.com/wiki/Cantina>, October 2019.

Seth Yalcin. *Modality and inquiry*. PhD thesis, MIT, 2008.

Seth Yalcin. Figure and ground in logical space. Posted at <http://escholarship.org/uc/item/11c0x4n5>, on 12–02–2015, 2015.

Seth Yalcin. Belief as question-sensitive. *Philosophy and Phenomenological Research*, 2016. ISSN 1933-1592. doi: 10.1111/phpr.12330. URL <http://dx.doi.org/10.1111/phpr.12330>.