

# Figure and Ground in Logical Space\*

Seth Yalcin  
yalcin@berkeley.edu

April 22, 2011

## 1 Introduction

The idea that states of belief are, in a certain sense, sensitive to *questions*, or to *subject matters*, or more generally to *ways of resolving logical space*, helps in some simple ways with aspects of the classical problem of logical omniscience. So I argue here. Focusing on belief, I begin by reviewing a version of a familiar story about belief and belief content, what I will call the *map picture* of belief. I will suggest that the picture is incomplete in ways that lead to the problems of logical omniscience, and that the addition of the aforementioned kind of sensitivity helps to fill in the picture in ways that start to address the problems. My larger aim is to explore the extent to which the idea of belief as question-sensitive state can be motivated by considerations in the philosophy of content, considered largely in abstraction from issues in descriptive semantics *per se* (e.g., in abstraction from the detailed compositional semantics of belief ascription). By the end, we will not have fully resolved the problems of logical omniscience, but we will have made some headway.

## 2 The map picture

The motto of the map picture is: *belief is the map by which we steer*.<sup>1</sup> You will have heard some version this story before, but we need a single version of it for operating on. So let me review the main features of this picture as I want to understand it. The picture owes substantially to [Stalnaker \[1984\]](#).

---

\*Note added 11/28/14: This paper was presented on a conference on alternatives in semantics held at Cornell University in April 2011 organized by Will Starr and Matti Eklund. It works up material in the third chapter of my dissertation ([Yalcin \[2008\]](#)). A number of the basic ideas here also appear in [Yalcin \[2011\]](#). After I wrote this paper, I set it aside and forgot about it until I noticed that it somehow got cited by [Yablo \[2014\]](#). I am greatly indebted to Yablo for the discussions that ultimately lead to this paper. Aside from the addition of some references, notably to Alejandro Pérez Carballo's illuminating related work ([Pérez Carballo \[2014\]](#)), and the occasional passage reworded for clarity, I have left the paper as-was.

<sup>1</sup>The motto was popularized by Armstrong (see, e.g., [Armstrong \[1973\]](#)), and goes back to Ramsey.

A belief state is a state of an agent that represents the world as being a certain way (the map aspect), and which plays a certain role in the explanation of the action (the steering aspect). The representational dimension of belief we understand in terms of the idea of information.<sup>2</sup> Crudely, a state carries information about the environment when it systematically co-varies in the appropriate way with the condition of the environment, so that the condition of the state serves as an indicator how things are in the environment. We take belief to be a state which tends, under normal conditions, to carry information in this sense; it is a state that tends, normally, to indicate how things in the world are. The representational (informational) content of a state of belief we model as a set of possible worlds, the set of possibilities that are the way the state represents the world to be. More precisely, we embrace:

**The possible worlds model of content.** The content of a state of belief is representable as a set of metaphysically possible worlds, intuitively the worlds “left open” by what is believed or accepted. Propositions are sets of possible worlds, and the propositions an agent accepts are those true with respect to all of those worlds the state leaves open.

(Not that the use of the map metaphor somehow entails a possible worlds model of content. Again, I am merely stipulating that this is part of what I will want to mean by “the map picture”.) What makes some set of possibilities an agent’s *belief* content is the role that these possibilities play in explaining how the agent ‘steers’. To be in a belief state whose content is a set of worlds  $B$  is in important part to be disposed to act in ways that would tend to satisfy one’s desires, were one in an  $B$ -world. We could think of it like this: given an agent’s desires and behavior, we solve for the content of their belief by finding the class of possibilities where their behavior is prone to lead to the satisfaction of their desires. Or better, we can see this approach as solving for belief and desire content simultaneously: one’s belief content is the belief component of that belief-desire package which best rationalizes one’s actions, and which otherwise maximizes the extent to which the belief state is an information-carrying state. As it is sometimes put, the picture takes it that belief-desire content is “constitutively rational”.<sup>3</sup>

---

<sup>2</sup>Some will prefer to understand the representational dimension, not in terms of causal or information-theoretic connections, but rather in terms of eligibility (see Lewis [1983, 1986]).

<sup>3</sup>Classic attempts to spell out “systematically co-varies in the appropriate way” and “tends, normally” include Dretske [1981, 1988], Stalnaker [1984], Fodor [1990a,b]. For an attempt to clarify the interdependence of the contents of belief and of desire, see Dretske [1988]. For more discussion of the idea of constitutive rationality, see Lewis [1994], Stalnaker [2004].

Belief has a holistic character on this picture. One solves for the content of an entire belief state at once, not one belief at a time. The individual beliefs an agent has are understood derivatively, as the propositions true throughout their unstructured belief content. Thus to have a single nontrivial belief is to have a flood of other beliefs. We have nothing close to an informative correlation between the individual believed-propositions and the complexity of the agent's internal state; there is, especially, no internal belief box containing a mentalese sentence for each proposition believed-true. Rather than the metaphor of an internal list of sentences, what we have instead is the map metaphor. States of belief represent richly, like maps. If a map represents things so that Oakland is east of San Francisco and south of Berkeley, it also represents everything entailed by Oakland's being east of San Francisco and south of Berkeley. For instance, it represents things so that San Francisco is west of Oakland, Berkeley is northeast of San Francisco, San Francisco and Berkeley are nonidentical and spatially apart, and so on. All these entailments are already included by the information given by the map. No extra ink needed. Consequences are free. In this way, the possible worlds model of content is maplike.<sup>4</sup>

A map is worth a thousand propositions, but in a sense, it is indeterminate which thousand. A map represents a way the world might be, but it does not do so *via* representing a set of propositions, as, say, a list of sentences does. The information in a map might be distilled into a set of propositions, but this can be carried out in myriad ways, and the map itself does not—or at least, it doesn't obviously—privilege one way over any of the others. If there is any proposition which has a privileged status according to a map-like representation, it is the single proposition which states: *things in the world are arranged thusly*. That is, it is the proposition which is just the content of the map.

Suffice to say the picture just sketched is controversial at every turn. I of course have no plan to fully defend the picture here, but only to address, in a limited way, some of the problems it faces. But to motivate that project a bit more, let us briefly review some old and familiar reasons for thinking that something is right about the map picture.

First, it is constitutive of this model that beliefs do not come and go one at a time, and that belief change is holistic in nature. That seems to correspond with reality. When you form the belief that you left your keys in the car, you thereby learn that the keys are not in your pocket, and that they are not in this room, and that they are not in the building, and that to get your keys, you will have to go your car, and so on. When the propositions we believe change, they

---

<sup>4</sup>Strictly, to call consequences 'free' is a bit misleading, insofar as this suggests a special class of propositions antecedently 'paid for'. The point is that, of themselves, entailments of what is believed don't impose any additional representational burden over their entailers.

change as a whole system, and in a way that preserves their overall coherence. (Compare the way that the information conveyed by a map changes, when we shift the location of an item on it.)

Second, we are often quite content attributing beliefs to agents whether or not the proposition said to be believed is one they have ever actively considered. Take Clyde, who swerves to avoid hitting a stray moose on the highway. It seems safe—i.e., true—to say that Clyde believes that the moose he just barely missed is larger than a golfball. Pragmatically odd, but safe. This, even though no thoughts of golfballs have recently crossed Clyde’s mind. Our comfort with the ascription seems related to the fact that, had Clyde not believed this, he wouldn’t have swerved so hard. The belief ascription seems appropriate because classifying the content of his state of belief with respect to this proposition does work explaining, *inter alia*, Clyde’s failure to treat the moose as golfball-sized.

Once we let in one example like this, it is easy to see how to multiply them. The result is that if we want to vindicate the appearances of ordinary belief-desire explanations, it seems we have a need to allow for a whole body of implicit belief. A vast body of implicit belief is predicted, and made intelligible, by the possible worlds map picture. Ordinary belief-desire explanation needn’t require the imputation of implausible level of explicit cogitation on the part of the agent being explained.

Thanks to this, the possible worlds map picture facilitates explanations of communication and of coordinated behavior by appeal to the ideas of common belief and common knowledge, concepts whose application characteristically entails the truth of belief and knowledge ascriptions of unbounded complexity. On the map picture, we needn’t impute mental representations of unbounded complexity to agents who are said to have common belief or common knowledge.

And then there are non-human agents. We ordinarily explain the behavior of non-human animals via the attribution of beliefs and desires, in a fashion entirely analogous to the explanation of human behavior. The map picture of belief can accept such explanations at face value. Attributions of belief on this picture do not involve the ascriber in any commitments about the specific structure of mental representation; *a fortiori*, there is no risk of representing ascribers as misconstruing this structure when they characterize the beliefs of agents of radically different cognitive design.

So the map picture, as articulated with the possible worlds model of content, has some compelling features. Plausibly it is the holistic character of belief on this story that accounts for the strengths just mentioned. But plausibly that same holism accounts for its central and best-known problem, a problem widely thought to vitiate the whole account: the problem of logical omniscience.

### 3 Problems of deduction

Actually, a variety of problems have been placed under the heading “the problem of logical omniscience”. We can start by separating out some of the problematic features of the possible worlds model.

First:

**(N) Closure under necessary equivalence.**

If  $A$  bears the belief relation to  $p$ , and  $p$  and  $q$  are true in exactly the same worlds, then  $A$  bears the belief relation to  $q$ .

This property appears to license intuitively implausible belief ascriptions— notably, those brought to attention by Frege’s puzzle about identity statements, and those resulting from the conflation of every necessary proposition. This is of course a very serious difficulty, but it is one I set aside for the purposes of this paper. Aspects of this problem arguably results less from the possible world’s model of belief content than from its model of propositions;<sup>5</sup> and arguably the resources needed for dealing with this problem are different in character than those required to for dealing with the other problems of deduction afflicting the possible world model.<sup>6</sup>

What other problems? I have in mind those arising as a result of the following three closure properties:

**(C) Closure under conjunction.**

If  $A$  bears the belief relation to  $p$ , and  $A$  bears the belief relation to  $q$ ,  $A$  bears the belief relation to the proposition which is the conjunction of  $p$  and  $q$  (namely,  $p \cap q$ ).

**(M) Closure under believed material implication.**

If  $A$  bears the belief relation to  $p$ , and  $A$  bears the belief relation to  $p \supset q$ , then  $A$  bears the belief relation to  $q$ .

**(E) Closure under entailment.**

If  $A$  bears the belief relation to  $p$ , and  $p$  entails  $q$  (i.e.,  $p$  is true at  $w$  only if  $q$  is true at  $w$ ),  $A$  bears the belief relation to  $q$ .

---

<sup>5</sup>For instance, a Montague-Scott model of belief, according to which a belief state is represented not as a set of worlds but as a sets of sets of worlds (with this set construed as the set of believed propositions), would be equally afflicted by (N).

<sup>6</sup>See however [Pérez Carballo \[2014\]](#) for an account of mathematical beliefs very much in the spirit of the present paper.

It will be agreed that these are all highly counterintuitive as necessary properties of belief states. Here we see that the map picture, as articulated with the possible worlds model, suffers an unfortunate irony: taking the belief-desire states of agents to be constitutively rational, it makes no room for the intelligibility of deductive reasoning as a rational activity. It is hard, on this picture, to see even what deductive reasoning would be. Whatever it is, it is not the sort of thing a rational agent would do, since rational agents already believe all of the consequences of their beliefs.

Intuitive this isn't. We should like to revise the possible worlds model of content so as to escape these three closure properties. What is the best of way doing this, while preserving as much possible the advantages of the map picture?

## 4 Atlases and fragmentation

One popular approach revises the possible worlds model of content as follows. Rather than representing the beliefs of an agent by a single set of possibilities, as on the standard model, we represent the beliefs of an agent by a collection of sets of possibilities. On this enriched picture, agents do not have a single belief state; rather they have a set of belief states, or “seperate systems of belief”, the contents of which are each represented by a set of worlds. These systems of belief may be “compartmentalized” from each other, so that the content of one system may include  $p$  but not  $q$ , and another  $q$  but not  $p$ . In such a case the agent may fail to believe some consequence  $r$  of  $p$  and  $q$ , because  $p$  and  $q$  are believed with respect to distinct belief states.<sup>7</sup> Braddon-Mitchell and Jackson [2007] offer an illustration, in the context of advocating this kind of theory:

Jones may believe that Mary lives in New York, that Fred lives in Boston, and that Boston is north of New York, but fail to put all this together and form the belief that Mary will have to travel north to visit Fred. ... [However] Jones may, consistently with the theory, have a system of belief according to which  $P$  and a *different* system of belief according to which  $Q$ , and so fail to believe that  $P\&Q$  by virtue of not having a system of belief according to which  $P\&Q$ . Indeed, it makes good sense that subjects should have different systems of belief, just as travellers often have a number of maps that they use on their travels. (199)

And Stalnaker writes:

... it may be a *nontrivial problem to put separate beliefs together into a single coherent system of belief*. All of my actions may be rational in that

---

<sup>7</sup>Stalnaker [1984] seems to have been the first to suggest this kind of approach. See also Lewis [1988a].

they are directed toward desired ends and guided by coherent conceptions of the way things are even if *there is no single conception of the way things are that guides them all*. There may be propositions which I would believe if I put together my separate systems of belief, but which, as things stand, hold in none of them. These are the propositions whose truth might be discovered by a purely deductive inquiry.

[Stalnaker, 1984, (54)] (emphasis mine)

So adjust the guiding metaphor. Rather than a single map, an agent's beliefs determine something more like a set of maps. Each map in the set is internally consistent, but it may be that some of the maps conflict with each other about how things are. We still steer by a map at any given time, but not always by the same map. We could try putting the new motto like this: *belief is the possibly inconsistent atlas from which we select maps by which we steer*. (Less catchy, alas, but more accurate, or so it is suggested.)

Note that since agents are now allowed to have a number of states or systems of belief, believing is actually a *three*-place relation between an agent, one of the agent's belief states, and a proposition: one believes a proposition *with respect to* some state *s* of belief.<sup>8</sup> Since we are multiplying the number of belief states one can have, it may help to introduce a term for the state of having the set of belief states one has. Let me use *doxastic state* in this way. So on this representation, an agent has in a single doxastic state (the atlas), and this consists in having a multiplicity of belief states (the maps). The doxastic state, we could say, is fragmented in that the belief states it consists in are compartmentalized from each other. I will refer to this as *the fragmented possible worlds model of content*, and I will refer to its intuitive gloss as *the atlas picture*.

**Fragmented possible worlds model of content.** The content of a state of belief is representable as a set of metaphysically possible worlds, intuitively the worlds 'left open' by what is believed. Propositions are sets of possible worlds, and the propositions an agent believes *with respect to a belief state* are those true with respect to all of those worlds the state leaves open.

On the fragmented model, (C) and (M) are still true, but *only if one holds fixed which state one is talking about*. If one does not, (C) and (M) are not really

---

<sup>8</sup>There is alternative way of putting the proposal, depending on how we want to use 'agent'. If we adopt an idealized conception of agency according to which agents are definitionally understood to have only a single state of belief, then we can keep the familiar idea that believing is a two-place relation between an agent and a proposition. What we lose, on this use, is the assumption that we are agents. Rather, beings like us are represented as a bundle, or system, of agents. I will return to this way of thinking of things briefly in section 7.

well-stated, because they advert to a two-place, rather than three-place, belief relation. What we can say is that the following modified closure properties fail on the fragmented model:

**(C') Closure under conjunction.**

If  $A$  bears the belief relation to  $p$  with respect to some state, and  $A$  bears the belief relation to  $q$  with respect to some state,  $A$  bears the belief relation to the proposition which is the conjunction of  $p$  and  $q$  (namely,  $p \cap q$ ) with respect to some state.

**(M') Closure under believed material implication.**

If  $A$  bears the belief relation to  $p$  with respect to some state, and  $A$  bears the belief relation to  $p \supset q$  with respect to some state, then  $A$  bears the belief relation  $q$  with respect to some state.

The idea of fragmentation enables a local, context-sensitive kind of holism about belief, one which preserves the various advantages of the map picture enumerated above. It does add complexity to the representation of belief, but this kind of additional complexity looks inevitable, at least if one wants to keep a representation of belief content broadly along these lines but also allow for the possibility of inconsistency.

I will take it without further discussion that fragmentation is indeed the best way to evade the counterintuitive consequences of (C) and (M). My chief concern in this paper is with the trouble arising from (E), closure under entailment. Given that we now assume a fragmented model, let us restate (E) in the obvious way:

**(E') Closure under entailment.**

If  $A$  bears the belief relation to  $p$  (with respect to some belief state  $s$ ), and  $p$  entails  $q$  (i.e.,  $p$  is true at  $w$  only if  $q$  is true at  $w$ ),  $A$  bears the belief relation to  $q$  (with respect to  $s$ ).

Here we face what is arguably the heart of the problem of deduction as it afflicts possible worlds accounts of content. The move to fragmentation does nothing to help us with manifest implausibility of closure under entailment.

Again, we require some revision to our model. And again, we should like a revision which preserves as much as possible the advantages of the map—or rather, atlas—picture.

En route to finding an appropriate adjustment to the fragmented model, we might pause to inquire into etiology of fragmentation. Why exactly is compartmentalization a fact of life for realistic agents? Understanding the nature



of this limitation is important if we are to really understand the atlas picture and the idea of compartmentalized belief. Advocates of fragmentation typically advert in some way or other to the computational limitations of the agent. It is suggested that agents do not have anything close to the representational resources needed to register, or make available for cognition and action, all of the consequences of their beliefs. Without this capacity, it is inevitable that inconsistencies will go largely unregistered, and fragmentation unchecked.

This kind of reply is intuitive, and I think on the right track, but it should be observed that it employs distinctions that are not yet integrated into the abstract framework being assumed (the fragmented possible worlds model). What is for a proposition to be “available” in the relevant sense, and in virtue of what does this state of mind impose some nontrivial computational cost? It is not clear how these ideas are supposed to be taken on the picture as so far described. Computational costs are the sort of thing incurred by the need to process some representations; but so far, our picture is completely neutral concerning the question how exactly belief content is encoded or represented. (We have employed the metaphor of a map, but of course the account does not require that mental representations literally be maps.) To be in a belief state whose content is a set of worlds  $s$ , for all we have said, is just to (i) be in a certain kind of indicator state, and (ii) to be disposed to act in ways that would tend to satisfy one’s desires, were one in an  $s$ -world. This approach to settling the content of a state of belief does not make or require any distinction, among what is believed, between a special class of “cognitively available” propositions and the rest.

It would be a mistake to forsake the account of implicit belief that the atlas picture provides. But if want an account of the content of belief that goes beyond what is implicit, and which can mark a distinction between cognitively accessible propositions and the rest, I suggest we extend the fragmented picture enough to include the idea of accessibility.

## 5 Resolutions of logical space

If we start, as the map picture does, with a picture of belief according to which it is primarily an attitude to possible states of the world and not to a class of propositions distinguishing those states, how can we make room for a special class of ‘available propositions’ without fundamentally distorting the picture? The response I want to consider begins with the recognition that, for realistic believers, the possible states that belief is an attitude towards are not maximally specific. They are coarse possibilities, possibilities reflecting answers to only so many questions. This coarseness betrays a trace of our computational

limitations, and of the notion of availability that we need. So a natural idea is that, insofar as we want a model of content that interfaces with these notions, information about this coarseness should somehow be included or reflected in the model. That is to say, the model should include information about the richness, or lack thereof, of the possible alternatives that an agent’s state of belief distinguishes, about what questions these alternatives speak to and fail to speak to.

What kind of formal element can play this role—can carry this kind of information in the model? As a first pass, we can consider the idea that a simple partition over logical space, a division of the space of possible worlds into mutually exclusive and jointly exhaustive regions, can play this role.<sup>9</sup> We can think of a state of belief as relativized to such a partition, and think of the cells of the partition as the non-maximal possible alternative situations ‘recognized’ by the state. So the thought is to move from a representation of a belief state that looks like (A) to one that looks like (B) (see fig. 1 below).

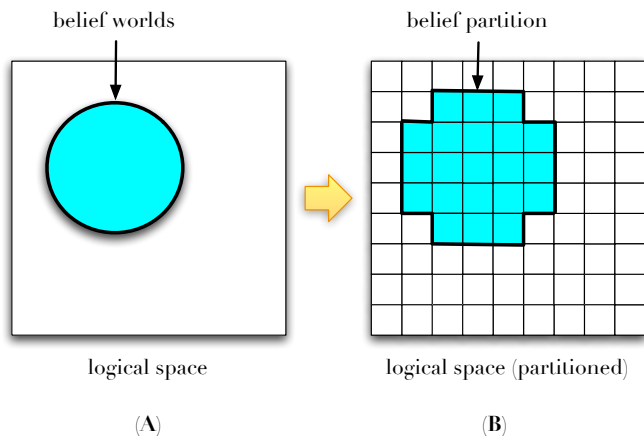


Figure 1: From possibilities as points to possibilities as partition cells.

A state of belief still determines a set of (maximally specific) possible worlds, but only insofar as it determines a coarser set of possibilities drawn from the partition of logical space in question. Call this latter set a *belief partition*. It partitions the (maximally specific) belief worlds. I will call a partition of logical space used in this role a *modal resolution*, or just a *resolution*. The cells of

<sup>9</sup>It should be noted that the idea that the context set is interestingly partitioned in a manner that is sensitive to the distinctions of interest to the interlocutors is an idea Stalnaker has raised in various places; see, e.g., Stalnaker [1981, 1986, 2014].

the resolution represent ways the world might be, but they are not maximally specific in the manner typically assumed of possible worlds; they settle some but not all questions. A resolution, we can say, *foregrounds* some distinctions and *backgrounds* others. With respect to a resolution, the *foreground* propositions are those propositions whose truth value supervenes on the resolution. (Equivalently: those propositions constructible entirely from unions of cells of the resolution.) We could say that the resolution makes these propositions visible, or brings them into focus. The rest of the propositions are, with respect to the resolution, *background* propositions.

Now the idea is that states of belief are relativized to resolutions. As before, the *believed* propositions correspond to those which are true throughout the belief worlds; but among the believed propositions we may now distinguish, in the obvious way, *foreground* beliefs and *background* beliefs. That enables us to consider the following modeling proposal: the propositions which constitute the available information of a body of belief content are those beliefs which are foregrounded by the resolution. Background beliefs correspond to the ‘unavailable’ information implicit in the foreground beliefs. Figure 2 depicts the distinction.

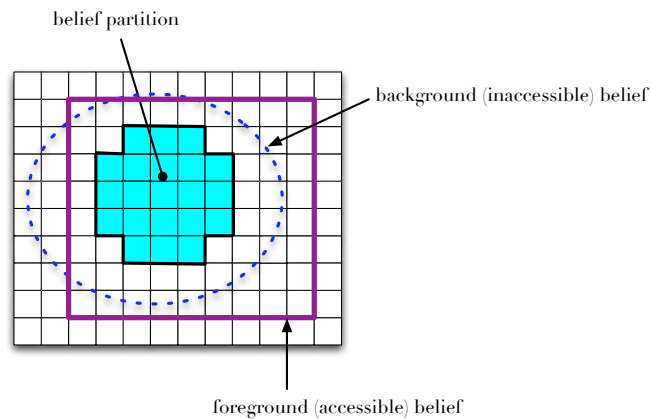


Figure 2: The distinction between available and implicit belief content.

On the fragmented possible worlds model, a doxastic state is represented as a set of sets of worlds. On the upgrade I now propose, a doxastic state is representable as a set of ⟨resolution, subpartition⟩-pairs. Actually, I will additionally assume that, given any particular resolution of logical space, a doxastic state maps it to at most one belief partition. Thus we can, somewhat

more intuitively, model a doxastic state as a (partial) function from resolutions of logical space to a set of cells taken from that resolution. Think of it like this: given a menu of (coarse, mutually exclusive, jointly exhaustive) alternatives for how things might be, as represented by a resolution, a doxastic state selects some set of these alternatives as its candidates for the way things actually are.

We can put the idea another way. Suppose we follow Hamblin [1958] (cf. Belnap and Steel [1976], Groenendijk and Stokhof [1984]) in modeling *questions* as partitions of logical space. On this account, the individual partition cells correspond to the possible complete answers to the question, and the various possible unions of partition cells reflect the possible answers. Then the current proposal can be stated like this: a doxastic state is a function from questions to answers. The answers may only be partial, leaving more than one coarse alternative open. And the question reflected by a resolution needn't be one particular easy to express in language. It may be preferable to understand it as a capturing a family of topically related questions on which the doxastic state takes a stance—as capturing a relatively detailed project of inquiry. We can say a belief state is *sensitive* to a question just in case the question is included in some question the state is defined upon. (One question is *included* in another just in case every cell of the latter is a union of cells of the former.) We take it belief states are not, or any rate need not be, sensitive to every possible question. There are ever so many questions I have not in any sense framed. If one is out to model realistic agents—as we are—it will be sensible to restrict to doxastic states defined on a finite number of resolutions, and to restrict to resolutions making only finitely many distinctions.<sup>10</sup>

There is yet another way to put the idea. A *subject matter*, according to Lewis [1988b,a], is partition of logical space, one that forms equivalence classes of worlds depending on whether the worlds yield the same verdict concerning the subject matter. (Cf. the notion of an ‘issue’ in Hulstijn [1997].) It divides up logical space, but only so far as concerns those distinctions native to the subject matter. To use Lewis’s example, if the subject matter is demography, then two worlds demographically alike in all respects will fall into the same cell within the partition determined by the subject matter, though they may differ in any other respect. We say that a proposition is *about* a subject matter just in case the truth-value of the proposition supervenes on the subject matter; equivalently,

---

<sup>10</sup>A number dynamic semantic systems for questions (e.g., Hulstijn [1997], Groenendijk [1999]) effectively assume that the common ground of a conversation should be modeled as a partitioned set of worlds, with the partition reflecting the question which is live in conversation. If the common ground of a conversation is supposed to reflect the mental states of presupposition of the interlocutors (as in, e.g., Stalnaker [1978]), then these theorists are, in my terminology, proposing that the state of presupposition is a resolution-sensitive state. (I agree.)

if the proposition is identical to some unions of cells from the partition. If one has this representation of subject matters, then my proposal can be stated like this: a doxastic state is a function from subject matters to some information about that subject matter.

Summing up:

**Resolution-sensitive model.** A state of belief is representable as a partial function mapping a resolution of logical space (question, subject matter) to a belief partition (answer, information about the subject matter). An agent’s accessible beliefs, relative to a resolution, will be those propositions true throughout the belief worlds and foregrounded by the resolution. Those propositions true throughout the belief worlds but backgrounded by the resolution are the agent’s inaccessible, implicit beliefs.

Observe that wherever you have a finite space of alternatives of the sort presented by a resolution, you can give a measure of how much information—in bits—would be needed to reduce those alternatives to one. This is just the logarithm, to the base 2, of the number of alternatives.<sup>11</sup> The toy resolution in (B) above, for instance, distinguishes 81 possibilities; so given only this representation of the alternatives and no other information, it would require at least 6.33985 bits of information to cut those possibilities down to one. I will call this number in bits the *information potential* of a resolution. To get an intuitive grip on what this number means, look at a resolution from the perspective of this question: *what is the smallest number of propositions needed to construct it?* The answer to this question will be the information potential of a resolution, rounded up to the nearest integer. If a resolution consists of only four cells, its information potential is 2 bits, since one minimally needs two propositions to effect a four-way division of logical space. More generally, if the cardinality of a resolution is equal to  $2^n$  for some  $n$ , then the number of propositions needed to draw the resolution will be  $n$ , and all of these propositions will be logically independent. If instead the size of a partition is  $m$  and  $2^n < m < 2^{n+1}$ , you will need  $n$  logically independent propositions *plus* one additional proposition entailed by one or more of those  $n$  propositions to construct the resolution.

The point of flagging the concept of the information potential of a resolution is to show that the model makes space for the idea of a clear, finite upper

---

<sup>11</sup>Hartley [1928] was the first fix on a logarithmic measure, and the first to attempt to measure information in terms of the number of possibilities eliminated.

On a Bayesian extension of the current picture, the alternatives are allowed to vary in probability. On such a model we might prefer to use the Shannon entropy of the relevant belief state as a measure of its informational content (Shannon [1948]).

bound on how much information—at least in one very specific, technical sense of ‘information’—a given belief state might in principle carry. This is an important point, and it is a point at which our abstract story of content, of what is represented, makes contact with what does the representing. Although the idea of resolution-sensitivity does not come with any very specific commitments about the detailed form by which the content of belief is represented, it does at least demand that the form have the complexity to *encode*, in bits, the information potential of the resolution in question. Resolutions impose what we could call an *encoding cost*. This is a way, albeit an abstract way, that this framework for representing belief imposes a constraint on a model of the form or mechanism of representation.

## 6 Accessible beliefs and their inaccessible consequences

Back to our last problem of deduction, and our last problematic closure property. Reformulate this property in the obvious way, in order to make it compatible with the assumptions of our new, resolution-sensitive version of the fragmented possible worlds model.

**(E'') Closure under entailment.**

If  $A$  bears the belief relation to  $p$  (with respect to some resolution  $\Pi$ ), and  $p$  entails  $q$ ,  $A$  bears the belief relation to  $q$  (with respect to  $\Pi$ ).

Now, this property does indeed hold on the model. But it applies in generality only to implicit, background beliefs. If one restricts to the class of propositions foregrounded by the belief state—the propositions relevantly accessible, or available—we observe this class is not closed under entailment. If an accessible belief  $p$  entails  $q$ , it does not follow that  $q$  is an accessible belief. Most of the consequences of one’s accessible beliefs are inaccessible.

I call this (modest) progress. To give the abstract model some more intuitive content, it will help to consider some examples of failures of deductive omniscience, and observe how the current framework deals, or can deal, with them. Consider the following two examples (from [Stalnaker \[1984\]](#)):

- William III of England believed that England could avoid war with France. Did he believe that England could avoid nuclear war with France? We want to say no; or at least, we want to be able to explain what is off about this characterization of his mental state.

- The absent-minded detective believes the butler did it, but totally overlooks the possibility it was the chauffeur. Does the detective believe that the chauffeur did not do it? We want to say no; or at least, we want to be able to explain what is off about this characterization of his mental state.

Taking each of these in turn:

William III had, of course, no thoughts about matters nuclear *per se*. Whatever concepts are, it seems natural to say that he lacks the concept NUCLEAR. Within the present framework, the modeling idea will be that his doxastic state is not defined on any resolution of logical space sensitive to questions of nuclear war. Again, it nowhere carves logical space according to whether nuclear war *per se* will, or will not, occur.

The absent-minded detective, by contrast, suffers from no particular lack of conceptual resources. He is capable of being sensitive to the relevant possibilities. He just isn't. Like a beginning chess player, his problem is that he fails to organize the alternative possibilities according to the distinctions that will facilitate the achievement of his aims. His state of mind is defined on the question: *Is it the butler or not?* But it is undefined on the question: *Is it the butler, or the chauffeur, or someone else?* Again, the proposition believed is inaccessible and in the background.

It is admittedly not easy to say, in a very general way, what is it to be sensitive to a question in the relevant sense, in the sense we are formally modeling. It seems to be at least this: it is to be equipped with possible states that distinguish possible answers to the question, and it is to be receptive to information which speaks the question. (Our butler is deficient in the second respect; William III in both.) Understood in this way, question-sensitivity is the sort of thing simple measuring devices can manifest. My thermostat is equipped with possible that distinguish possible answers to the question, *within what range is the temperature in this room?*, and it is receptive to information which speaks that question. It is not equipped with possible states that distinguish possible answers to the question, *how is the weather in Topeka?*, and (a fortiori) is not receptive to information which speaks the question. What counts as being appropriately sensitive to information which speaks to a question is evidently a context-sensitive, comparative affair, at least as far as ordinary ascription goes.<sup>12</sup>

Beyond intuitive examples, a resolution-sensitive model allows us to say a somewhat more than the fragmented model about what is involved in deductive inquiry. Deductively thought, we can say, is generally a matter of dialing up the resolution of one's state or mind, a matter of bring consequences into the

---

<sup>12</sup>I discuss this issue a bit more in [Yalcin \[2011\]](#).

foreground; running with the metaphor, it is a matter of seeing consequences. We also get some perspective on what it is to bring compartmentalized states of mind together, and on why it is hard. Deductive thinking tends to require high resolution states of mind, states of mind with high information potential. Resolutions with high information potential have a high encoding cost. Although this is obviously not yet anything like a cognitive model of deduction, it does give us a sense of the connection between our failure to be deductive omniscient and our limited representational capacities. The high encoding cost of completely bringing together two compartmentalized states of belief—of bring the propositions accessibly believed with respect to each into the foreground at once—may be unaffordable for an agent with limited representational capacities.

## 7 Accessibility by degrees and local visibility

Our aim was to supply the atlas picture of belief, and its fragmentation model, with a distinction between accessible and inaccessible beliefs. We have done that. But note that our model is still idealized insofar as the class of accessibly believed propositions is closed under ‘visible’ logical consequence, in the following sense:

**(V) Closure under visible consequence.** If  $S$  accessibly believes that  $p$  (with respect to  $\Pi$ ), and  $p$  entails  $q$ , and  $q$  is foregrounded (with respect to  $\Pi$ ), then  $S$  accessibly believes  $q$ .

We should be happy to trade (E) for (V); still, is (V) an implausible constraint?

Consider a toy case where  $S$  accessibly believes, with respect to a given resolution, four logically independent propositions, and that let us represent this state of belief simply as a cell of a sixteen-cell resolution. Question: how many propositions is  $S$  represented as accessibly believing with respect to this resolution? According to the definition of available belief I have given, what is accessibly believed are those propositions true throughout the belief partition and also constructible from the cells of the resolution. How many propositions satisfy these criteria in this case? The answer is  $2^{15}$ , or 32,768. But surely that is totally implausible; surely that is more information than is actually accessible for the agent. Moreover, were we to assume that  $S$  believes some fifth logically independent proposition, again representing  $S$ ’s beliefs by a cell in the resulting resolution,  $S$ ’s available beliefs would number in the millions.

So there is a threat here that any resolution rich enough to characterize a nontrivial body of belief is also one which will make for *far too much* available belief. This seems to me less a reason to abandon the resolution-sensitive



model—which, despite its commitment to (V), is an order of magnitude less abstract and idealized than the standard possible worlds model from which we began—than to try to become clearer about just what it is for information to be accessible in the relevant sense. On this, the following remarks from Stalnaker may help set the mood:

There are questions I can answer quickly with a moment’s thought or a minor calculation, and questions that I have the computation resources to answer eventually, but only after a lot of time and effort. For some questions of the latter kind, I may be able to say outright that I have the capacity to produce the answer eventually; for others, I may in fact be able to produce an answer, if I choose the right computational strategy, but may be unable to say whether I can actually produce the answer. How easy must the search or computation be in order for the answer to count as something the agent already knows or believes, and not something it has the capacity to come to know or believe? ... There is obviously a continuum here, and no very natural place to draw the line between information that is easily accessible and information that is not. ([Stalnaker, 1991, 437])

My notion of availability may seem expansive, but just how expansive it really is depends on what assumptions are made about the average information potential of the resolutions that an agent’s doxastic state is defined on. On this I have not taken a stance. One wants a better grip on the notion of availability being modeled before one take a stance. Perhaps a realistic representation of a doxastic state on this model will be a function on very low resolutions of logical space—resolutions of, say, half a dozen cells or less, rather than the scores of cells that my diagrams have suggested. If so, the number of accessibly believed propositions at a given resolution will be reduced exponentially—rather than 32,000 available propositions, we will have 32 at most, assuming six cells. A doxastic state like this would presumably involve radical fragmentation.<sup>13</sup>

On this last point, let me note an obvious possibility. I have been assuming that a resolution of logical space either brings a proposition into focus—foregrounds it—or leaves it out of focus—backgrounds it. But there is a middle ground. Just as photograph might bring only part of an object into focus, so a resolution of logical space might make a proposition visible with respect only a region of logical space, leaving it backgrounded elsewhere. This would be to

---

<sup>13</sup> It would be natural to try to pair this kind of approach with attempts to represent intelligent thought and action as emergent from the interaction of a horde of primitive homunculi—see Dennett [1981], Fagin and Halpern [1987], Minsky [1988], and the introduction of Stalnaker [1999]—though note that the model I have described does not yet say anything about how these compartmentalized states of mind relate to one another.

allow for local visibility, visibility with respect to a partition of some proper subset of logical space.




Figure 3: Figure and ground in logical space: a resolution with a region in focus.

A simple illustration: in figure 3 we have a partition of logical space into four cells (as reflected by the dark lines). The proposition  $p$  is visible with respect to all of logical space. The proposition  $q$  is not, but it is visible with respect to the space of  $p$  worlds. The proposition  $r$  is invisible with respect both logical space and the space of  $p$  worlds, but is visible with respect to the  $p \& q$  worlds. One way of thinking of this is as a representation of an inquiry, or of a process of inquiry. An agent’s doxastic state may map the partition  $\{p, \neg p\}$  to  $p$ —that is, the agent may answer the question whether  $p$  affirmatively—and then ask, so to speak: ‘Given  $p$ , is  $q$  the case?’ in such a way that the agent may only be correctly represented as distinguishing  $q$ -visible alternatives within the  $p$ -space. Similarly for  $r$ . The representation of logical space here is zoomed in on the  $p$  worlds, and zoomed in again on the  $p \& q$  worlds, resolving this region more finely than any other area.

A resolution with local visibility like this can be thought of as an ordered sequence of simple resolutions over an increasingly small region of logical space, with each resolution thought of as a stage of a larger inquiry; and one can think of the sequence of resolutions itself as belonging to a coherent hierarchy of resolutions, one reflecting various paths through logical space the agent’s state of belief traverses. On this more complex model, whether or not a proposition counts as available would depend on one’s location in inquiry.

I note a second technical option for grading the notion of availability. Given a resolution, one can distill from it those smallest sets of propositions that are sufficient to determine the resolution, and treat these propositions in these sets as the most available propositions.<sup>14</sup> Easier, we could just pick one of these sets

<sup>14</sup>The cardinality of these sets will be the information potential of the resolution, rounded up to the nearest integer.

of propositions as being the most accessible. In either case, a gradable notion of availability could then be defined in terms of the complexity of formulae built out of propositional atoms expressing the most accessible propositions.<sup>15</sup> Alternatively, we might treat the propositions which are the cells of the resolution as the most accessible, again recovering a gradable notion in terms of complexity.

Whether any of these notions are interesting depends, obviously, on the application. If one is out to model the belief states of human beings, the current project shades into psychological questions of cognitive modeling and architecture, especially into questions of search and retrieval. Our aim has been less to make a first-order contribution to this project than to try to see it as unifiable with the representation of agents supplied by possible worlds accounts of content.

## References

- David Armstrong. *Belief, Truth and Knowledge*. Cambridge University Press, 1973.
- Nuel Belnap and Thomas Steel. *The Logic of Questions and Answers*. Yale University Press, New Haven, 1976.
- David Braddon-Mitchell and Frank Jackson. *The Philosophy of Mind and Cognition: An Introduction*. Blackwell Publishers, 2007. 2nd edition.
- Daniel C. Dennett. *Brainstorms: Philosophical Essays on Mind and Psychology*. MIT press, Cambridge, MA, 1981.
- Fred Dretske. *Knowledge and the Flow of Information*. MIT Press, Cambridge, MA, 1981.
- Fred Dretske. *Explaining Behavior: Reasons in a World of Causes*. MIT Press, Cambridge, MA, 1988.
- R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi. *Reasoning about Knowledge*. MIT Press, Cambridge, MA, 1995.
- Ronald Fagin and Joseph Y. Halpern. Belief, awareness, and limited reasoning. *Artificial Intelligence*, 34(1):39–76, 1987.
- Jerry Fodor. A theory of content, I: the problem. In *A Theory of Content and Other Essays*. MIT Press, Cambridge, MA, 1990a.
- Jerry Fodor. A theory of content, II: the theory. In *A Theory of Content and Other Essays*. MIT Press, Cambridge, MA, 1990b.
- Jeroen Groenendijk. The logic of interrogation. In *Proceedings of SALT IX*. Institute for Logic, Language and Computation (ILLC), University of Amsterdam, 1999.

---

<sup>15</sup>This bears some resemblance to a version of awareness logic. See [Fagin et al., 1995, sec. 9.5] for an overview.

- Jeroen Groenendijk and Martin Stokhof. *Studies on the Semantics of Questions and the Pragmatics of Answers*. PhD thesis, University of Amsterdam, 1984.
- Charles L. Hamblin. Questions. *Australasian Journal of Philosophy*, 36(3):159–68, 1958.
- Ralph Hartley. Transmission of information. *Bell System Technical Journal*, 7(3):535–563, 1928.
- Joris Hulstijn. Structured information states: raising and resolving issues. In A. Benz and G. Jäger, editors, *Proceedings of MunDial97*, University of Munich, 1997.
- David K. Lewis. New work for a theory of universals. *Australasian Journal of Philosophy*, 61(4):343–377, 1983.
- David K. Lewis. *On The Plurality of Worlds*. Blackwell, Malden, MA, 1986.
- David K. Lewis. Relevant implication. *Theoria*, 54(3):161–174, 1988a.
- David K. Lewis. Statements partly about observation. *Philosophical Papers*, 17(1):1–31, 1988b.
- David K. Lewis. Reduction of mind. In S. Guttenplan, editor, *A Companion to the Philosophy of Mind*, pages 412–431. Blackwell, Oxford, 1994.
- Marvin Minsky. *Society of Mind*. Simon and Schuster, New York, 1988.
- Alejandro Pérez Carballo. Structuring logical space. *Philosophy and Phenomenological Research*, Online first version(DOI: 10.1111/phpr.12116):1–32, 2014.
- Claude E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27(3):379–423, 1948.
- Robert Stalnaker. Assertion. In Peter Cole, editor, *Syntax and Semantics 9: Pragmatics*, pages 315–332. Academic Press, 1978.
- Robert Stalnaker. Indexical belief. *Synthese*, 49(1):129–151, 1981.
- Robert Stalnaker. *Inquiry*. MIT Press, Cambridge, MA, 1984.
- Robert Stalnaker. Possible worlds and situations. *Journal of Philosophical Logic*, 15(1):109–123, 1986.
- Robert Stalnaker. The problem of logical omniscience, I. *Synthese*, 89(3):425–440, 1991.
- Robert Stalnaker. *Context and Content: Essays on Intentionality in Speech and Thought*. Oxford University Press, Oxford, 1999.
- Robert Stalnaker. Lewis on intentionality. *Australasian Journal of Philosophy*, 81(1):119–212, 2004.
- Robert Stalnaker. *Context*. Oxford University Press, Oxford, 2014.
- Stephen Yablo. *Aboutness*. Princeton University Press, 2014.
- Seth Yalcin. *Modality and Inquiry*. PhD thesis, Massachusetts Institute of Technology, Cambridge, MA, 2008.
- Seth Yalcin. Nonfactualism about epistemic modality. In Andy Egan and Brian Weatherson, editors, *Epistemic Modality*, pages 295–332. Oxford University Press, Oxford, 2011.