

Belief as Question-Sensitive

SETH YALCIN

University of California, Berkeley

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 small ways with aspects of the problem of logical omniscience. Also it tentatively suggests a certain way of thinking about what it is to possess a concept. With extensions, it suggests an approach to modeling disagreement about which questions are relatively more “natural”, and about which questions are “fully factual”.

So I suggest here. I begin by reviewing a version of the familiar possible worlds model of belief and belief content. The picture is incomplete in ways that lead to the problem of logical omniscience. I will suggest that the addition of the aforementioned kind of sensitivity helps to fill in the picture in ways that start to address the problem. I then describe some (optional) applications of this picture of belief for modeling concepts and for modeling inquiry into, and disagreement about, which distinctions are natural or structure-tracking, and about which distinctions are “fully factual” or real. 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 natural language semantics about belief ascription. By the end we certainly will not have fully resolved the problems of logical omniscience, but we will have made some headway.

2. Belief as Map

Start with the metaphor of belief states as maps. The motto of the map picture is: *belief is the map by which we steer*.¹ You will have heard some version this idea before. I am interested specifically in a certain the way this picture of belief can be developed following in the path of Hintikka (1962), using tools from possible worlds semantics. The version I give owes substantially to Stalnaker (1984).²

¹ The motto was popularized by Armstrong (see, e.g., Armstrong 1973), and goes back to Ramsey (1931).

² In this paper I broach the idea of question-sensitivity (or subject matter-sensitivity, or resolution of logical space-sensitivity) for belief specifically from the perspective of possible worlds approaches to modeling belief content: I will present question-sensitivity as offering solutions for some of the problems this view faces. But theorists who favor structured conceptions of content still have reason to read on, since as I hope to show, question-sensitivity can do some of the work that structured content is often thought to do. The issue of question-sensitivity interacts with the debate about structure.

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 agent's actions (the steering aspect). The representational dimension of belief we understand in terms of the idea of information.³ 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 of 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. Propositions are sets of possible worlds, and the propositions an agent believes 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.) 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 a *B*-world. More fully, 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 can be characterized as an information-carrying state. As it is sometimes put, the picture takes it that belief-desire content is constitutively rational.⁴ “One predicts behavior in such a case by ascribing to the system the possession of certain information and by supposing it to be directed by certain goals, and then by working out the most reasonable or appropriate action on the basis of these ascriptions and suppositions” (Dennett 1971: 90). This is in the spirit of the conception of agency one finds in decision theory—but of course the possible worlds model is even more idealized, modeling a belief state with just a set of worlds rather than a probability distribution over the space of worlds.

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 propositions believed 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

³ 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. For a critical discussion of Lewis's view, see Stalnaker 2004).

⁴ Classic attempts to spell out ‘systematically co-varies in the appropriate way’ and ‘tends, normally’ include Dretske (1981a, 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).

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 map-like.

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 might. 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 me briefly review some of the basic attractions of the possible worlds version of 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 also come to believe that the car is where your keys are, that the keys are in a vehicle, that the keys are not in your pocket, that they are not in the room you are in, that to get your keys, you will have to go your car, that some keys are in some car, and so on. These are not discrete cognitive achievements. When the propositions we believe change, they change as a whole system, and in a way that tends to preserve 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 comfortable attributing beliefs to agents whether or not the proposition said to be believed is one they have ever actively considered. Suppose you swerve to avoid hitting a stray moose on the highway. It seems true to say that you believe that the moose you just barely missed is larger than a golfball. This, even though no thoughts of golfballs need recently have crossed your mind. Our comfort with the ascription seems related to the fact that, had you not believed this, you wouldn't have swerved like that. The belief ascription seems appropriate because classifying the content of your state of belief with respect to this proposition does work explaining (*inter alia*) your failure to treat the moose as golfball-sized.

It is easy to see how to multiply examples like this. Such examples suggest that if we want to vindicate the appearances of many ordinary belief-desire explanations, 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 relative to 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 nontrivial difficulty. What I will have to say in this paper bears only indirectly on this problem, which arises owing specifically to the conception of propositions deployed in the possible worlds model of content, and not especially to the closure assumptions built into it.⁵ Arguably some of 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 worlds model. (See however Pérez Carballo 2014 for an account of mathematical beliefs very much in the spirit of the present paper.) I will say just a little bit about this issue, in section 7.

Turn meanwhile to the other problems. I have in mind those arising as a result of the following three closure properties:

⁵ For comparison, consider a Montague-Scott model of belief, according to which a belief state is represented not as a set of worlds, but instead as a set of sets of worlds (with this set construed as the set of believed propositions). A Montague-Scott model can easily avoid assumptions of logical omniscience—the set of believed propositions needn't be closed under entailment—but it will have the same problems with (N) as does the Hintikka-style possible worlds model.

(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 .

These are all obviously problematic as necessary properties of belief states. The map picture, as developed in possible worlds style, suffers an unfortunate irony: taking the belief-desire states of agents to be constitutively rational, it makes little 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 believer would do, since such believers are modeled as already believing all of the consequences of their beliefs.

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 standard 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 “separate 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. Stalnaker (1984) seems to have been the first to suggest this kind of approach; see also Lewis (1988a). 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 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 realistic, perhaps.)

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 of belief.⁶ 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 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. A doxastic state is representable as a set of belief states. The content of a belief state is representable as a set of metaphysically possible worlds, intuitively the worlds "left open" by what is believed relative to the state. 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 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:

(C²) Closure under conjunction (fragmented).

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.

(M²) Closure under believed material implication (fragmented).

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.

⁶ There is alternative way of putting the proposal. If we adopt an idealized conception of agency according to which agents are definitionally understood to have only to 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 10.

The idea of fragmentation enables a local, context-sensitive kind of holism about belief, one which preserves much of what is advantageous about the map picture. It does add nontrivial complexity to our formal 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 take it that fragmentation goes some distance in addressing counterintuitive consequences of (C) and (M). So I take fragmentation on board. My chief concern in this paper is with the trouble arising from (E), closure under entailment. As with (C) and (M), (E) remains true if one holds fixed which state one is talking about. But even if one does not explicitly hold fixed which state one is talking about, closure under entailment holds. That is, the following modified closure property holds:

(E²) Closure under entailment (fragmented).

If A bears the belief relation to p with respect to some state, and p entails q , A bears the belief relation to q with respect to some state.

Thus the move to fragmentation does nothing to ease the pain caused by 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 about etiology of fragmentation. Why exactly should compartmentalization be a fact of life for realistic agents? Understanding the nature of this limitation is important for understanding how to conceive of fragmentation. Advocates of the fragmentation picture 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 (cf. Stalnaker (1991, 1999b)). 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 modeling framework being assumed (the fragmented possible worlds model). What is for a proposition to be “available” or “accessible” 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—and indeed, that was supposed to be a virtue of the approach. (We have employed the metaphor of a map, but of course the account does not require that mental representations literally be maps.) We are thinking of belief as a kind of indicator state that stands in rationalizing functional interconnections to states of desire and to action. This conception of belief, and the corresponding conception of what kind of facts fix the content of a state of belief, does not obviously make or require any distinction, among what is believed, between a special class of “cognitively available” propositions and the rest.

We shouldn't like to lose the account of implicit belief that the atlas picture provides. But if we want an account of the content of belief that goes beyond what is implicit, which can mark a distinction between cognitively accessible or available propositions and the rest, we might try reaching for a model which can make explicit some distinction between accessible and inaccessible content.

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.⁷ 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 initial, very simple, 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).

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 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 that cut according to the resolution: they are 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.

⁷ See Stalnaker (1981, 1986, 2014) for intimations construable in something like this direction. The need for "coarse possibilities" also recommends taking seriously versions of modal logic wherein possibilities needn't fix the truth values of all formulas, as in possibility semantics (Humberstone 1981, Rumfitt 2015, Holliday 2015b). See sections 8 and 9 below.

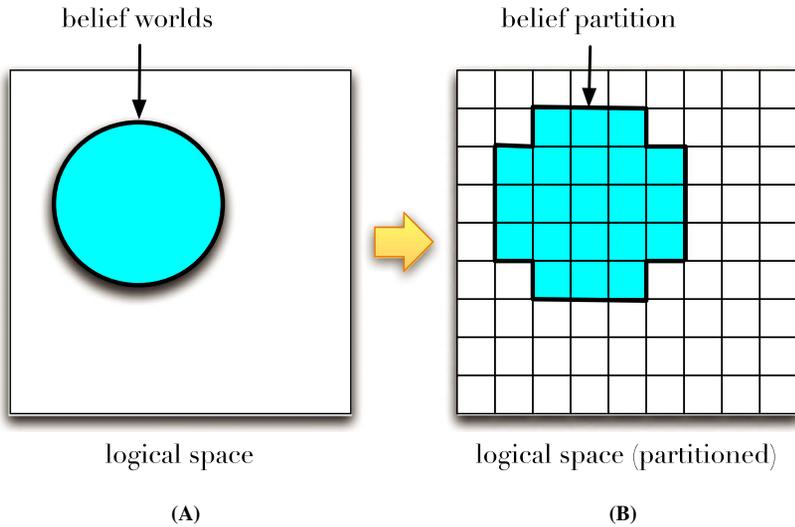


Figure 1. From possibilities as points to possibilities as partition cells.

Now the idea is that states of belief are relativized to resolutions of logical space. 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, using two example propositions.

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

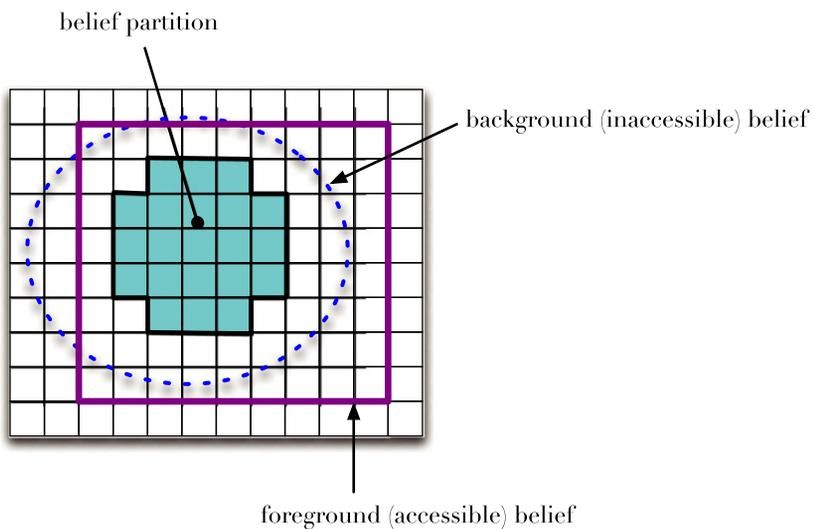


Figure 2. The distinction between available and implicit belief content.

of pairs of a resolution and a subpartition at that resolution. 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 (for where the actual world is to be found).

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 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 former is a union of cells of the latter.) 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, 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.⁸

There is yet another way to put the idea. A *subject matter*, according to Lewis (1988a, b), can be modeled as a 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 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. (Cf. Yablo 2014.)

Here is one last kind of spin. Various theorists have, for various reasons, entertained the idea of thinking of a state of belief as a triadic relation between a subject, a proposition, and a mode of presentation of that proposition (see Schiffer 1977, 1992, Crimmins and Perry 1989, Crimmins 1992). The present account can be construed as a sort of

⁸ A number of dynamic semantic/pragmatic systems for questions (e.g., Hulstijn (1997), Groenendijk (1999); cf. Roberts (1996, 2012)) 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 (cf. Stalnaker (1978)), then these theorists are, in my terminology, proposing that the state of presupposition is a resolution-sensitive state. (I agree; see Yalcin (2011).)

possible worlds version of this idea. Given a doxastic state and a resolution it is defined on, we could think of that resolution as supplying something like a mode of presentation for the propositions it makes visible. The cognitive import of a proposition turns partly on the matrix of background distinctions it is considered relative to.

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.⁹ 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. 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 note that the model makes space for the idea of a clear, finite upper bound on how much information—in one specific technical sense of ‘information’—a given belief state might in principle carry. This is a point where our abstract story about content, about what is represented, makes nontrivial 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 a very abstract way, that this framework for representing belief imposes a constraint on a model of the form or mechanism of representation. (Of course, the notion of information in play here is only remotely connected to more intuitive notions of information.)

⁹ Hartley (1928) was the first to 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. (For discussion see Yalcin 2007, 2011.) 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).

6. Accessible Beliefs and their Inaccessible Consequences

Back to our last problem of deduction, and our last problematic closure property. In the resolution-sensitive version of the fragmented possible worlds model, we can separate two relevant closure properties:

(E³) Closure under entailment (resolution-sensitive).

If A bears the belief relation to p with respect to some resolution Π , and p entails q , A bears the belief relation to q with respect to some resolution Π' .

(E⁴) Closure under entailment (with visibility).

If A bears the belief relation to p with respect to some resolution Π where p is visible, and p entails q , A bears the belief relation to q with respect to some resolution Π' where q is visible.

It is not difficult to see that (E³) holds while (E⁴) fails. If one restricts to the class of propositions foregrounded by the belief state—the propositions relevantly accessible, or available—this class is not closed under entailment. The failure of (E⁴) helps to draw out the point that most of the consequences of one's accessible beliefs are inaccessible.

I suggest this is some kind of progress, speaking to at least part of what troubles us about closure under entailment. 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. Were he to consider this possibility, it might shake up his view. 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 (see the next section), 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 relevant lack of ability to cut logical space. He is capable of being sensitive to the relevant possibilities. He just isn't. Somewhat 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?* The possibility overlooked is inaccessible and backgrounded.

(Actually there is a second possible route of response. Confronted with the question *Is it the butler, or the chauffeur, or someone else?*, the detective might have a clear disposition to answer: *either the butler or the chauffeur*. So it may be right, on some versions of this case, to describe his doxastic state as defined on that question. The issue is then that he isn't disposed to steer by that map in his atlas save in highly limited or particular circumstances—e.g., the circumstance of another person saying to him “What about the chauffeur?”, or seeing the chauffeur twirl his mustache, etc.)

To be sensitive to a question seems at least to be equipped with possible states that distinguish possible answers to the question, and 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 states 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?*; it is not receptive to information which speaks that question.¹⁰

Deductive thought, we can try saying, is generally a matter of dialing the resolution of one's state of mind into one that renders the premises and conclusion visible. (Not to suggest that availability implies conscious availability; it does not.) Maybe what we call deductive reasoning is nontrivial in part because it tends to require higher resolution states of mind, states of mind with high information potential. Resolutions with high information potential have a high encoding cost. This idea would connect our failure to be deductive omniscient to our limited representational capacities. The high encoding cost of completely bringing together two compartmentalized states of belief—of bringing the propositions accessibly believed with respect to each into the foreground at all once—may sometimes be unaffordable for agents like us.

7. Relocating Concepts

Let us return briefly to (N), to closure under necessary equivalence. On the face of it, nothing we have done makes any advance on the problems (N) gives rise to, since if *A* believes *p* relative to some resolution, and *p* and *q* are true in exactly the same worlds, then *A* believes *q* relative to that resolution. Further, necessities remain indistinguishable and inescapable: all belief-having creatures seem to come out believing things like (1) and (2), including my cat:

- (1) Bachelors are unmarried.
- (2) Brothers are male siblings.

Confronted with this very basic problem, many philosophers will be inclined to reach for finer-grained objects to play the proposition role—in particular, for a structured conception of propositions. On some versions of the structured approach, concepts are key

¹⁰ I discuss the issue of what it is to be sensitive to a question a bit further in Yalcin (2011).

to the story. Consider the Fregean tradition in philosophy of thinking of propositions as built from concepts (or senses, or modes of presentation). My cat doesn't believe what (1) says, on this approach, partly because he doesn't have thoughts built from concepts like *BACHELOR* and *UNMARRIED*. He doesn't possess these concepts. One's possession of a concept is reflected in one's capacity to entertain thoughts in which the concept figures, and my cat lacks the relevant capacities.

The present picture suggests a different way of understanding what a concept is, and of understanding what it takes to possess one. To possess a concept is to have an ability to cut logical space in a certain way, to distinguish possibilities in terms of the sorts of things that answer to the concept. Concepts can be construed formally as much like subject matters in Lewis's sense—or workably limited fragments thereof, corresponding to questions that “demand not the whole truth, but only, say, a paragraph-length answer that hits the highlights” (Lewis 1988a: 162). A concept determines a matrix of distinctions. For example, the concept/subject matter *BACHELOR* corresponds to the partition of logical space distinguishing possibilities depending on what's happening with the bachelors at each world—so that two worlds will belong to the same cell just in case they don't differ in their bachelor respects. To possess a concept, on this idea, is to be capable of entering states of mind sensitive to the associated distinctions.¹¹

This puts concepts in a different place in our cognitive economy. Concepts are not literally parts of the propositions we believe, though the propositions we (availably) believe are in a sense cut from their cloth. One's deployment of concepts is reflected, not in the proposition one believes, but rather in the resolution of logical space it is believed relative to. As Yablo (2014) suggests, the aboutness of a mental state resides not just in the proposition one is related to, but also in the matrix of distinctions—the subject matter—it is located within.¹²

We recoil from the claim that my cat believes bachelors are unmarried, and from things like it. This might partly be explained by appeal to the thought (in play already for the William III example) that such ascriptions often do communicate something about the conceptual resources of the subject—about the kind of resolutions the subject's state of mind is defined on. This might be especially so when the propositional content ascribed is necessary, for such ascriptions can't be explanatory of behavior in the paradigm way that assignments of content are supposed to explain behavior in the possible worlds setting.

More ambitiously, subject matters might be deployed to disentangle intensionally equivalent concepts—say, the concept *WATER* and the concept *H₂O*, if you think those are intensionally equivalent. Start with the point that subject matters can contain (include) other subject matters. Examples: “A big subject matter, the 17th Century, includes the smaller, more specialized subject matter, the 1680's. A big subject matter, how many stars there are, includes the smaller subject matter, whether there are finitely or infinitely many stars” (Lewis 1988a: 162). It seems natural to say that even if ‘water’ and ‘H₂O’ have the same intensions/pick out the same property, which in turn fixes the same

¹¹ Pérez Carballo (2014) pursues related lines of thinking; I have benefitted from his instructive discussion. See also Stalnaker (1998).

¹² See Yablo (2014)'s notion of *directed content*—what he models as a pair of a proposition and a subject matter. Yablo anticipates my emphasis on subject matter: one of his central ideas is exactly that “subject matters are an important and potentially independent factor in meaning, over and above truth-conditional content” (21).

partition of logical space, the concepts WATER and H₂O are nevertheless distinct because they correspond to embeddings (inclusions) of that partition into different subject matters.¹³ For example, the former concept might (*inter alia*) be understood as embedding the partition into a subject matter reflecting parochial human interests and concerns, one including, say, the subject matter *beverages*—so that with the concept WATER, we (*inter alia*) locate that stuff amongst the beverages—while the latter embeds it in some part of the subject matter *chemistry*—so that with the concept H₂O, we *inter alia* locate that stuff amongst the chemicals.

This is all programmatic and tentative; much more would have to be said to turn this into a theory. None of it should be taken to suggest that the issues raised by (N) have been fully resolved. A fully satisfying account here would require (*inter alia*) a semantics and pragmatics for belief ascription, one which shows how exactly the expressions used in a *that*-clause influence our assumptions about what resolutions the subject's state of mind is defined on. It remains to relate concepts, understood in this kind of way, to linguistic meaning.¹⁴

In the next two sections, I describe some optional applications and extensions of the preceding for modeling certain kinds of inquiry: first, inquiry into questions of objective structure (“nature’s joints”), and second, inquiry into which distinctions are “fully factual” in character.

8. Natural Questions

Finding the truth is epistemically good. Often finding the truth is a matter of asking the right question. So asking the right question is epistemically good. Bringing question-like structure into our model of belief can help us to make sense of what it could mean to ask the right question.

Some questions—some ways of carving up logical space—are more natural than others. “Are emeralds green?” is a more natural question than “Are emeralds grue?”, which is in turn a more natural question than “Are emeralds grue, and where can I get a taco at this hour?” We can say that the naturalness of a question turns on whether it involves relatively more natural properties—‘natural’ in the sense of Lewis (1983, 1986). Or we can put it as Sider (2011) does, in terms of objective structure. A partition of logical space can do a better or worse job at carving at the joints—at separating possibilities in a manner that hews to the objective structure of reality. An idea suggests itself: one key good-making feature of a question is naturalness. Better to be in a doxastic state defined on relatively more natural questions.

¹³ On this line of thinking, concepts could be modeled as partitions with a distinguished subpartition, the latter corresponding to the associated intension.

¹⁴ This topic is broached, inconclusively, in Yalcin (2011). This incidentally points to one way my approach in this paper differs from the (congenial) project of Schaffer (2007), who argues that knowledge is question-sensitive largely from considerations about the semantics of knowledge ascription. (Most of his considerations would not carry over directly to ‘believes’, which does not have the same potential to embed questions as ‘knows’.)

Schaffer draws on, and contributes to, the “relevant alternatives” tradition in thinking about knowledge (Dretske 1981b). The connections between relevant alternatives epistemology and the present approach to belief remain to be mapped. Theorists in this tradition want some version of the idea that to know a proposition, one must eliminate certain alternatives; and it seems clear that the space of alternatives must form a partition on this view (see Holliday 2015a for discussion).

Asking the right question has instrumental epistemic value, through its role in finding truth. Does its epistemic value run deeper? Sider (2011) writes:

The goal of inquiry is not merely to believe truly (or to know). Achieving the goal of inquiry requires that one's belief state reflect the world, which in addition to lack of error requires one to think of the world *in its terms*, to carve the world at its joints.

Joint-carving thought does not have merely instrumental value. It is rather a constitutive aim of the practice of forming beliefs, as constitutive as the more commonly recognized aim of truth. (61)

It seems a nontrivial question of epistemology whether Sider is right that joint-carving thought is an epistemic end in itself. But to frame and pursue this question, we should like to clarify the underlying state of mind. What is it to think of the world in its terms? One picture of this presupposes that thought has a language-like structure. "Thinking of the world in its terms" is thinking propositions constructed from the ideology privileged by the world ('ideology' used in the sense going back to Quine 1951)—for instance, it is thinking propositions built from joint-carving concepts (or structured by the content of such concepts). Sider seems to have something like this picture, one he signals on his first page:

The world has a distinguished structure, a privileged description. For a representation to be fully successful, truth is not enough; the representation must also use the right concepts, so that its conceptual structure matches reality's structure. There is an objectively correct way to "write the book of the world". (i)

I only want to note that we can agree with much of the spirit of this without an especially linguistic picture of mentation. The representational structure Sider takes to be answerable to nature's joints might be located, not in the propositions believed, but in the way that the logical space one believes relative to is carved.¹⁵

All this relates to substantivity (Sider 2011: ch. 4). Some questions are more substantive than others. The question "Do electrons repel?" is more substantive than "Is Robinson Crusoe a bachelor?" Sider proposes that the substantivity of an issue turns on whether it concerns more or less joint-carving distinctions. As we are putting it, nonsubstantive questions are less natural than substantive ones. Suppose so. The question arises where to locate views about substantivity within our model of believing. We might agree about the answer to a question, and yet differ about whether the question is substantive. What it is to disagree in this way? It need not be disagreement about which world is actual. Maybe we both agree that "some nonsense made out of sour green apple liqueur" served in a cocktail glass is not a martini (Bennett 2009), but our agreement might be shallow in the following sense: I regard the distinction between martinis and nonmartinis as basically terminological, whereas you, the committed mixologist, see the distinction as pretty close to a joint in nature. This might be so even if we associate 'martini' with the same possible worlds intension. To get at the difference between our attitudes about the substantivity of the distinction, we can appeal to the idea that a question-sensitive state of belief includes also a *naturalness ordering*, a partial order on the questions the state

¹⁵ Questions arise here about how this shift in perspective interacts with Sider's proposal to move "beyond the predicate" and extend the notion of structure to logical notions. I hope to broach this issue elsewhere.

of mind is defined on. Two believers may be alike in what questions their states of mind are defined on and alike in respect of their answers to these questions, but differ in respect how they judge the relative naturalness of the questions. Strictly ideological differences between agents can be modeled in this way.

I just said that disagreement about what distinctions are substantive need not be disagreement about which world is actual. This does not imply that such disagreements have no basis in reality, that they are not about the world. We can hold that there is an objectively privileged preorder of questions by naturalness. (Indeed we must hold this, if we want to maintain the just-suggested picture of what it is to disagree about substantivity while also concurring with Sider that the matter of which questions are substantive is objective.) One kind of epistemic achievement is to endorse only correct answers to questions; another is to have the right view about which questions are natural.¹⁶

What is the objectively privileged naturalness ordering like? (For example, does it have a top element—a maximally natural way of distinguishing possibilities? Does the best “map of the world” occur at a specific resolution?) This is a question of metaphysics. The present framework doesn’t take a stand; rather, it enables us to model what it is to agree or disagree about answers to this question.¹⁷

9. Inquiry into the Reality of Structure

Disagreement about whether a certain distinction is substantive is separate from disagreement about whether that distinction is, as we could put it, *real*. Perhaps you think the bachelor/nonbachelor distinction is not very substantive. Still, you are apt to think that certain things are bachelors, and others not—that there are various bachelor facts, that one can mark real distinctions between possibilities according to what is going on with their respective bachelors. Nonsubstantive questions typically have factual answers. It is just that their answers aren’t especially revelatory of the structure of the world.

Disagreement about the reality of a distinction is another thing. Consider those debates in metaphysics characterized by an opposition between *relationist* (or *structuralist*, or *comparativist*) views on the one hand and what we could call *absolutist* views on the other hand. These debates arise in metaphysical discussions of things like spatial locations and distances, the direction of time, mass, simultaneity, and orientation. To a first approximation, absolutists about a domain ascribe strictly more objective structure to the domain—in the sense of recognizing more distinctions within it—than their relationist counterparts. Further, they typically offer to ground facts that a relationist would take as primitive in that additional structure. Example: a prototypical relationist about space will hold (roughly) that there are no facts about the spatiotemporal location of a material thing that go beyond the pattern of distance relations between it and other material things. A

¹⁶ Against the vein of this thinking, one might hold instead that the natural/structural distinctions are fixed once we fix a world. Possibilities which are “categorically” alike might nevertheless differ in their “facts of naturalness”. If it is contingent which questions are natural, then the joint-carving ways of dividing logical space will vary from world to world. This would imply that to know what world one is in is also to know which distinctions are structural; and that disagreement about which questions are substantive is, at bottom, disagreement about what world one is in. I don’t have the space to argue against this view. I only have the space to note that it is not mandatory—not the only way of vindicating the idea that some questions are more substantive, more objectively natural and structure-tracking, than others.

¹⁷ See Yablo (2014) for additional applications of subject matter-sensitivity to problems of epistemology, including classic problems in confirmation theory.

prototypical absolutist about space, by contrast, will recognize more such facts, namely facts about the location of the thing within space itself. These latter facts are, furthermore, the fundamental facts of location for material bodies, says the absolutist; the facts of relative location for material bodies are derivative from these. Another example: a relationist about mass will hold that there are no facts about the mass of an object x that go beyond the facts of comparative mass—the fact that x weighs twice as much as object y , half as much as object z , and so on. The absolutist about mass accepts all these relative mass facts, but recognizes also a class of facts of absolute mass. Further, she takes the latter facts to be more fundamental. She will ground the fact that x is half the weight of z in the absolute mass facts about x and z —as it might be, in the fact that x has the absolute monadic mass property of being one kilogram and z has the absolute monadic mass property of being two kilograms.¹⁸

From the vantage point of the present picture of belief, we can model the disagreement between a relationist and an absolutist as a difference about which ways of cutting up logical space are real. But here, a nontrivial adjustment to our model is in order. We are well-advised to take on a conception of logical space broader than the picture of it as a set of maximally specific ways that a world might metaphysically be. For we need room in the picture for (not merely unnatural, nonsubstantive distinctions but) unreal distinctions. A suitable upgrade to our picture of logical space might be delivered using the tools of possibility semantics (Humberstone 1981, Holliday 2015b). Think of possibilities now as fundamentally region-like rather than point-like.¹⁹ Finer and finer possibilities settle more and more, and are located relative to richer and richer sets of distinctions. Some of these distinctions are the metaphysically real distinctions; others are “merely conceptual” distinctions, projected artifacts of our models, distinctions that don’t correspond to real differences.²⁰ On this conception of logical space, the metaphysical possibilities are not the points of the space; rather, they can be construed as the elements of a privileged way of dividing the space, one where possibilities are separated only by real distinctions.²¹

Then what we can say is that theorists who disagree about whether a distinction is real disagree, in a sense, about “how specific reality is”—about which ways of cutting logical space correspond to factual differences, and which correspond to something more like mere artifacts of our representations, to projected structure, to distinctions that are not

¹⁸ See Dasgupta (2013, 2015) for overviews of these debates.

¹⁹ So possibilities may overlap, and are akin to sets of points. But perhaps the space is completely point-free; or perhaps some regions but not others can be refined to points. (Structure might be postulated without end; there may be no limit to inquiry.)

²⁰ Note: while this amounts to recognizing possibilities that are not metaphysically possible, I am not here dropping the requirement that possibilities be consistent. (Some possibilities go in for more structure than there really is, but there is no possibility where there is a round square cupola on Berkeley College.)

²¹ We might say that the metaphysical possibilities form a partition of logical space in the following sense: given a set W of possibilities with a partial order \leq of refinement (the term used by Humberstone (1981)), define a partition of the poset $\langle W, \leq \rangle$ to be a maximal set of pairwise incompatible possibilities, where x and y are incompatible iff they have no common refinement according to \leq . So while an ordinary partition of a set X is a subset of the powerset of X , a partition in the sense just defined is simply a subset of X . (I am indebted to Wes Holliday for suggesting this extended notion of ‘partition’.)

Strictly speaking, in the possibility-theoretic setting, one might deny that the space of metaphysical possibilities forms a partition even in this sense. This turns partly on the question whether some metaphysical possibilities are refinements of (are more “specific” or “determinate” than) others. I am abstracting away from this issue here.

fully factual. The absolutist will favor a finer grain than the relationist, seeing a certain distinction in logical space as factual and real where the relationist will see only a non-factual or artifactual distinction. Actuality as it is according to the absolutist about (e.g.) mass, on this picture, is a subregion of actuality as it is according to the comparativist.²²

On this way of modeling the disagreement, we needn't say that the absolutist about (say) mass disagrees with the relationist in respect of what 'mass' means. That would distort their debate, turning it quasi-semantic. The two sides can agree about which region of logical space corresponds to "The pineapple is two kilograms". What they disagree about is whether that region is visible at the grain of logical space favored by reality, the grain that foregrounds only factual distinctions. That is to say, they disagree about which logically possible distinctions are factual ones.

Where in the mind does one's view about the factuality of a question reside? A state of mind might be construed as bifurcating the set of questions is it defined on into those that are fully factual (according to that state of mind) and those that are less than fully factual (according to that state of mind). Then the difference between the absolutist and relationist comes to a difference in what propositions are visible according the questions they each regard as fully factual.

10. Accessibility by Degrees and Local Visibility

Our main aim was to supply the atlas picture of belief, and its fragmentation model, with structure for marking a distinction between accessible and inaccessible beliefs (while trying to stay true to the spirit of the picture). But note that our model is still highly idealized inasmuch 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 resolution Π , and p entails q , and q is visible at Π , then S accessibly believes q .

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

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. That seems like a lot—more information than we might think 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.

²² This approach has much in common with Jeff Russell's way of approaching debates about haecceity (Russell 2015). As Russell wants to put it, Lewis's "cheap haecceitism" is the view that haecceitistic propositions don't distinguish "factual possibilities"; rather, they are "thin propositions" that distinguish (not-fully-factual) possibilities. Russell offers to identify factual possibilities with equivalence class of "thin possibilities" (which, in the context of Lewis's cheap haecceism, are factual possibilities equipped with certain counterpart structure).

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 evolve it. On this, the following remarks from Stalnaker might set the tone:

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. A view on this should be guided by further examples and applications.²³ 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—with agents understood as constantly shifting the map they are steering by. If so, the number of accessibly believed propositions at a given resolution will be reduced exponentially. Doxastic states like this would involve radical fragmentation.²⁴

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 to a region of logical space, leaving it backgrounded elsewhere. This would be to allow for local visibility, visibility with respect to a partition of some proper subset of logical space.

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 $p \cap q$ 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 $\{p, \neg p\}$

²³ Natural language semantics/pragmatics may be one source of constraints. See Yalcin (2011), Cariani (2013), and recent work on inquisitive semantics.

²⁴ It would be interesting to compare 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 (1999a)—though note that the model I have described does not yet say anything about how these compartmentalized states of mind relate to one another.

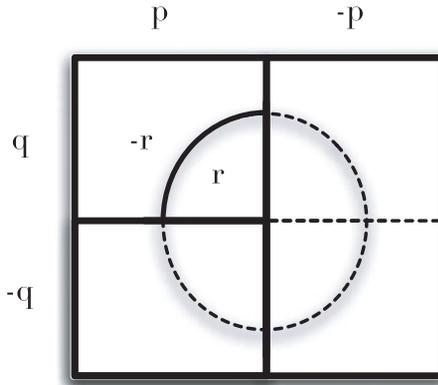


Figure 3. Seeing the independence of some c from a given p and q .

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 $p \cap q$ worlds, and zoomed in again on the $p \cap q$ worlds, resolving this region more finely than any other area.

Part of the idea here would be that as we progress in inquiry and resolve questions, the distinctions corresponding to those questions can fade into the background. To give a crude illustration, suppose you are thinking where best to vacation. First question: tropical locale or not? You decide tropical. Next question: luxury hotel or not? You decide luxury. Last question: summer or not? By the time you reach this question, ‘yes’ might seem obvious—you’re restricting now to luxury hotels in tropical places, and the salient alternative to summer will be the rainy season. Out of focus at this stage in your overall inquiry will be paradigm nontropical nonluxury nonsummer possibilities, like a cozy cabin in ski season. Maybe had you brought such possibilities into focus, it would have overturned your ultimate view about where best to vacation. But it seems that inquiry can, in the process of bringing new possibilities into focus, sometimes take turns that push other possibilities out of focus.

A resolution with local visibility like this might 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”. If propositions are understood as believed only relative to such inquiries, we can potentially make quite fine-grained distinctions between states of belief.²⁵

²⁵ 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. Easier, we could just pick one of these sets 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. This bears some resemblance to a version of awareness logic; see (Fagin *et al.* 1995: sec. 9.5) for an overview. 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.

An assumption of the current framework we should consider relaxing is the assumption that the alternative (full) answers to the questions our states of mind are defined on are mutually exclusive. Much fruitful thinking about questions does not operate with this assumption (Hamblin 1973, Karttunen 1977). I have erred on the Lewisian side of assuming partition-like structure, merely because it is good to see how far one can get under that restrictive assumption. But it would obviously be worthwhile to explore models relaxing it. Yablo (2014) already says this, and in particular explores the idea of moving from equivalence relations to similarity relations. Willer (2013)’s conception of epistemic possibility might suggest another direction for development.²⁶

Whether any of these notions are interesting depends, obviously, on the possibility of detailed applications. It may well be that the notion of availability divides into importantly different subtypes, that we have been blurring important distinctions. 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, retrieval, and framing (as pursued in, e.g., Kahneman *et al.* 1982, Gilovich *et al.* 2002).²⁷ Our aim has not so much been to make a first-order contribution to this project, as much as to try to see it as unifiable with the representation of agents supplied by possible worlds accounts of content.²⁸

References

- Armstrong, David. 1973. *Belief, Truth and Knowledge*. Cambridge University Press.
- Belnap, Nuel and Steel, Thomas. 1976. *The Logic of Questions and Answers*. Yale University Press.

²⁶ One place it may help to move beyond classic Lewisian subject matters, in connection with problems of logical omniscience, is in explaining the nontransparency of logical equivalences. Yablo (2014) raises the question: “If two statements— $P \wedge (Q \supset P)$ and $\neg(P \supset Q) \vee (P \wedge Q)$ —carve out the same region of logical space, how can the thinker fail to notice this?” (67). If we consider these sentences only relative to the simple four-way partition of logical space which groups worlds according to whether they agree on P and Q —that is, relative to the question “Which of P and Q are true?”—then we won’t have the resources to describe agents as differentially sensitive to them. But if we are allowed to move beyond partitions, the alternatives that a given proposition shares the foreground with may vary in interesting ways—and if we want, in ways that correspond to the structure of the sentence. Yablo’s example prompts the thought that the sentence $P \wedge (Q \supset P)$ picks out proposition relative to the distinctions in logical space marked by the elements of $\{P, Q, Q \supset P\}$ —its component pieces—while $\neg(P \supset Q) \vee (P \wedge Q)$ picks out the same proposition relative to a different, larger set of distinctions: $\{P, Q, P \wedge Q, P \supset Q, \neg(P \supset Q)\}$. On this picture, we abandon closure under visible consequence, and certain boundaries in logical space may get “traced over” multiple times.

²⁷ In a more empirical vein, there is good evidence that judgments of the probability of an outcome can be sensitive to the way that the alternatives to it are partitioned (Windschitl and Wells 1998; see Yalcin 2010 for discussion). If, as Hawthorne *et al.* (2016) suggest, what one counts as believing tracks what one views as likely, then since what one views as likely is a question-sensitive matter, so is what one believes. (See also Egan (2008) for another kind of empirical motivation for fragmentation.)

²⁸ This paper works up material in the third chapter of my dissertation (Yalcin 2008). A number of the basic ideas here also appear in Yalcin (2011) The first draft (which was entitled “Figure and Ground in Logical Space”) was presented at a conference on alternatives in semantics held at Cornell University in April of 2011 and organized by Will Starr and Matti Eklund. I am indebted to the audience at Cornell for helpful discussion. The most recent draft benefited from conversations with Wes Holliday, Niko Kolodny, Tamar Lando, Geoffrey Lee, Edward Schwartz, Robert Stalnaker, and from comments from an anonymous reviewer. I am especially indebted to Stephen Yablo, both for conversation and for the inspiration of his work. Yablo (2014) develops related themes in rich detail. Perez Carballo (2014) also overlaps in important ways with the present paper.

- Bennett, Karen. 2009. Composition, colocation, and metaontology. In David Chalmers, David Manley, and Ryan Wasserman, editors, *Metametaphysics*. Oxford University Press.
- Braddon-Mitchell, David and Jackson, Frank. 2007. *The Philosophy of Mind and Cognition: An Introduction*. Blackwell Publishers. 2nd edition.
- Cariani, Fabrizio. 2013. 'Ought' and resolution semantics. *Noûs*, 47(3): 534–558.
- Crimmins, Mark. 1992. *Talk About Beliefs*. The MIT Press.
- and Perry, John. 1989. The prince and the phone booth: Reporting puzzling beliefs. *The Journal of Philosophy*, 86(12): 685–711.
- Dasgupta, Shamik. 2013. Absolutism vs comparativism about quantity. In Karen Bennett and Dean W. Zimmerman, editors, *Oxford Studies in Metaphysics*, volume 8. Oxford University Press.
- 2015. Substantivalism vs relationalism about space in classical physics. *Philosophy Compass*, 10(9): 601–624.
- Dennett, Daniel. 1971. Intentional systems. *Journal of Philosophy*, 68(4): 87–106.
- 1981. *Brainstorms: Philosophical Essays on Mind and Psychology*. MIT Press.
- Dretske, Fred. 1981a. *Knowledge and the Flow of Information*. MIT Press.
- 1981b. The pragmatic dimension of knowledge. *Philosophical Studies*, 40: 363–78.
- 1988. *Explaining Behavior: Reasons in a World of Causes*. MIT Press.
- Egan, Andy. 2008. Seeing and believing: Perception, belief formation and the divided mind. *Philosophical Studies*, 140(1): 47–63.
- Fagin, Ronald and Halpern, Joseph. 1987. Belief, awareness, and limited reasoning. *Artificial Intelligence*, 34(1): 39–76.
- , —, Moses, Yoram, and Vardi, Moshe. 1995. *Reasoning about Knowledge*. MIT Press.
- Fodor, Jerry. 1990a. A theory of content, I: The problem. In his *A Theory of Content and Other Essays*. MIT Press.
- 1990b. A theory of content, II: the theory. In his *A Theory of Content and Other Essays*. MIT Press.
- Gilovich, Thomas, Griffin, Dale, and Kahneman, Daniel, editors, 2002. *Heuristics and Biases: The Psychology of Intuitive Judgement*. Cambridge University Press.
- Groenendijk, Jeroen. 1999. The logic of interrogation. In *Proceedings of SALT IX. Institute for Logic, Language and Computation (ILLC)*, University of Amsterdam.
- and Stokhof, Martin. 1984. *Studies on the Semantics of Questions and the Pragmatics of Answers*. PhD thesis, University of Amsterdam.
- Hamblin, Charles. 1958. Questions. *Australasian Journal of Philosophy*, 36(3): 159–68.
- 1973. Questions in Montague English. *Foundations of Language*, 10(1): 41–53.
- Hartley, Ralph. 1928. Transmission of information. *Bell System Technical Journal*, 7(3): 535–563.
- Hawthorne, John, Rothschild, Daniel, and Spectre, Levi. 2016. Belief is weak. *Philosophical Studies*, 173(5): 1393–1404.
- Hintikka, Jaakko. 1962. *Knowledge and Belief: An Introduction to the Logic of Two Notions*. Cornell University Press.
- Holliday, Wesley. 2015a. Fallibilism and multiple paths to knowledge. In Tamar Gendler and John Hawthorne, editors, *Oxford Studies in Epistemology*, volume 5.
- 2015b. Possibility frames and forcing for modal logic. UC Berkeley Working Papers in Logic and the Methodology of Science.

- Hulstijn, Joris. 1997. Structured information states: Raising and resolving issues. In Anton Benz and Gerhard Jäger, editors, *Proceedings of MunDial97*, University of Munich.
- Humberstone, Lloyd. 1981. From worlds to possibilities. *Journal of Philosophical Logic*, 10(3): 313–339.
- Kahneman, Daniel, Slovic, Paul, and Tversky, Amos, editors, 1982. *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge University Press.
- Karttunen, Lauri. 1977. Syntax and semantics of questions. *Linguistics and Philosophy*, 1: 607–653.
- Lewis, David. 1983. New work for a theory of universals. *Australasian Journal of Philosophy*, 61(4): 343–377.
- 1986. *On The Plurality of Worlds*. Blackwell.
- 1988a. Relevant implication. *Theoria*, 54(3): 161–174.
- 1988b. Statements partly about observation. *Philosophical Papers*, 17(1): 1–31.
- 1994. Reduction of mind. In Samuel Guttenplan, editor, 1994 *A Companion to the Philosophy of Mind*. Blackwell.
- Minsky, Marvin. 1988. *Society of Mind*. Simon and Schuster.
- Pérez Carballo, Alejandro. 2014. Structuring logical space. *Philosophy and Phenomenological Research*, Online first version (DOI: 10.1111/phpr.12116):1–32.
- Quine, W. V. O. 1951. Ontology and ideology. *Philosophical Studies*, 2(1): 11–15.
- Ramsey, Frank. 1931. General propositions and causality. In his *The Foundations of Mathematics and Other Logical Essays*. Routledge and Kegan Paul.
- Roberts, Craige. 1996. Information structure in discourse: Towards an integrated formal theory of pragmatics. In Jae-Hak Yoon and Andreas Kathol, editors, *OSU Working Papers in Linguistics*, volume 49. The Ohio State University Department of Linguistics.
- 2012. Information structure in discourse: Towards an integrated formal theory of pragmatics. *Semantics and Pragmatics*, 5(6): 1–69.
- Rumfitt, Ian. 2015. *The Boundary Stones of Thought: An Essay in the Philosophy of Logic*. Oxford University Press.
- Russell, Jeffrey. 2015. Possible worlds and the objective world. *Philosophy and Phenomenological Research*, 90(2): 389–422.
- Schaffer, Jonathan. 2007. Knowing the answer. *Philosophy and Phenomenological Research*, 75(2): 383–403.
- Schiffer, Stephen. 1977. Naming and knowing. *Midwest Studies in Philosophy*, 2(1): 28–41.
- 1992. Belief ascription. *The Journal of Philosophy*, 89(10): 499–521.
- Shannon, Claude. 1948. A mathematical theory of communication. *Bell System Technical Journal*, 27(3): 379–423.
- Sider, Theodore. 2011. *Writing the Book of the World*. Oxford University Press.
- Stalnaker, Robert. 1978. Assertion. In Peter Cole, editor, 1978 *Syntax and Semantics 9: Pragmatics*. Academic Press.
- 1981. Indexical belief. *Synthese*, 49(1): 129–151.
- 1984. *Inquiry*. MIT Press.
- 1986. Possible worlds and situations. *Journal of Philosophical Logic*, 15(1): 109–123.
- 1991. The problem of logical omniscience, I. *Synthese*, 89(3): 425–440.
- 1998. What might nonconceptual content be? *Philosophical Issues*, 9: 339–352.

- 1999a. *Context and Content: Essays on Intentionality in Speech and Thought*. Oxford University Press.
- 1999b. The problem of logical omniscience, II. In his *Context and Content*. Oxford University Press.
- 2004. Lewis on intentionality. *Australasian Journal of Philosophy*, 81(1): 119–212.
- 2014. *Context*. Oxford University Press.
- Willer, Malte. 2013. Dynamics of epistemic modality. *Philosophical Review*, 122(1): 45–92.
- Windschitl, Paul and Wells, Gary. 1998. The alternative-outcomes effect. *Journal of Personality and Social Psychology*, 75(6): 1411–1423.
- Yablo, Stephen. 2014. *Aboutness*. Princeton University Press.
- Yalcin, Seth. 2007. Epistemic modals. *Mind*, 116(464): 983–1026.
- 2008. *Modality and Inquiry*. PhD thesis, Massachusetts Institute of Technology.
- 2010. Probability operators. *Philosophy Compass*, 5(11): 916–37.
- 2011. Nonfactualism about epistemic modality. In Andy Egan and Brian Weatherson, editors, 2011 *Epistemic Modality*. Oxford University Press.

Addendum:

An error by *PPR* in the proof stage of the publication of this paper caused the following footnote to be left out:

A different implicit/explicit distinction is drawn within a fragmentation model of belief by Greco.¹ On Greco's way of theorizing the distinction, the idea is that one or more of an agent's belief fragments are of an explicit variety—these are relevant, especially, for what the agent is disposed to say in various situations—whereas other fragments are implicit: they control behavior in ways that often fly under the radar of the agent's articulated thinking. An example of Greco's: the “implicit sexist” John, who defends vigorously the equal intelligence of the sexes whenever given the opportunity, but who manifests, in his nonlinguistic behavior, the opinion that women are of lower intelligence (by paying attention to them less, by assigning their statements lower credence, etc.). An idea Greco considers, as I read him, is that John's explicit belief fragment is what is in control of his actions when he confronts questions like: “How should I talk about the intelligence of women?” Whereas most of the rest of the time, implicit fragments are steering the ship—for instance, when John is confronted with an issue like “How seriously should I take what Betty is saying?” or “How much should I pay attention to Betty?” (N.b.: by “confronting an issue”, I don't mean “subvocalizing an interrogative” — confronting a question is something nonlinguistic creatures can and do do, I take it.)

For Greco, some belief fragments are implicit and others are explicit, whereas for me, an implicit/explicit distinction (or an implicit/accessible distinction) is a distinction drawable for each individual belief fragment. But Greco's thinking seems complimentary to, and not incompatible with, what I am saying. I'd agree, for instance, that for his implicit sexist, we should talk about (at least) two fragments. I'd add that we could think of these fragments as each associated with different questions/subject matters, reflective of the sorts of situation the fragment is apt to control action. (And this seems to dovetail with aspects of Greco's informal talk of the purpose/task-relativity of belief and of knowledge.)

¹ Daniel Greco (2015) Iteration and Fragmentation. *Philosophy & Phenomenological Research* 91 (3) 656-73.