

Questions in Action

DANIEL.HOEK@VT.EDU, VIRGINIA TECH, JANUARY 2022.¹

ABSTRACT: Choices confront us with questions. How we act depends on our answers to those questions. So the way our beliefs guide our choices is not just a function of their informational content, but also depends systematically on the questions those beliefs address. This paper gives a precise account of the interplay between choices, questions and beliefs, and harnesses this account to obtain a principled approach to the problem of deduction. The result is a novel theory of belief-guided action that explains and predicts the decisions of agents who, like ourselves, fail to be logically omniscient: that is, of agents whose beliefs may not be deductively closed, or even consistent.

Imagine you are in a cold, dark forest at dusk, bereft of supplies and surrounded by disheartening animal noises. You come to a crossroads and have to choose a path. Hungry eyes are tracing you, and you face an almost palpable question: *How do I get out of here?* Questions await us at both the literal and metaphorical crossroads of life, even if they are usually less consequential. The choice of how many eggs to buy raises the question *How many eggs go into a carbonara for four?* Plotting your next chess move, you face the question *How do I put my opponent on the defensive?* In the flower shop, you wonder *What is his favourite colour?* And so on. Whenever you make a choice, you face a question.

Moreover, what you decide to do normally depends on your *answer* to the question the choice raises. Take the supermarket situation. If you reckon you need five eggs for the carbonara, you will buy half a dozen. If you think you need eight, you get a dozen. If you are unsure, maybe you still get a dozen just to be on the safe side. Thus your choice is guided by your answer to the question it confronted you with. If you know the correct answer, you will generally choose well (or achieve your immediate aims, anyway), while wrong answers lead to bad choices. When faced with a question you cannot answer, you are likely also unsure what to do. Our beliefs, then, are the answers we have to the questions that our choices confront us with.

This paper develops this question-centric or *inquisitive* way of thinking about the relation between beliefs and choices. It is divided into two parts. First, Part A gives a precise articulation of an inquisitive belief-action principle, and shows how this leads to a more refined understanding of belief-guided action. Part B then uses this new belief-action principle as the basis for a principled approach to a long-standing problem in the theory of belief: the problem of deduction, also known as the problem of logical omniscience. This is the difficulty of saying how deductive inquiry can be fruitful, given that, by its nature, deduction gives us no new information (e.g. Dummett 1975).

¹ Winner of the Isaac Levi Prize; to appear in *The Journal of Philosophy*. The ideas in this paper have been cooking for quite a while, benefiting from more interactions than I can acknowledge here. My greatest debt is to my doctoral advisor, Cian Dorr. A two-hour meeting with Cian is a uniquely intense and invigorating philosophical experience. It took a great many of those to get this project into shape and see it through. For insightful comments on drafts of the paper, I am grateful to Bob Beddor, Sam Carter, Ivano Ciardelli, Kevin Dorst, Paul Egré, Adam Elga, Cody Fenwick, Jane Friedman, Simon Goldstein, Ben Holguín, Nico Kirk-Giannini, Harvey Lederman, Matt Mandelkern, Salvador Mascarenhas, Agustín Rayo, and especially Dave Chalmers and Jordan MacKenzie. Parts of the paper were written during stays at the Institut Jean Nicod and Princeton, and the work was supported with funding from NYU's Global Research Initiatives and Virginia Tech's Kellogg Center for PPE.

The inquisitive belief-action principle developed below is based on the classical belief-action principle on which Ramsey (1926, 1927) built his decision theory. The latter makes no reference to questions, stating simply that a given belief manifests as a general disposition to act on the informational content of the belief. But consider this puzzle case, inspired by Powers 1978 and Elga and Rayo 2021a:

ROMEO RECALL: Juliet comes home to find a note saying “*Somebody called for you — didn’t catch a name but he sounded upset.*” There is a phone number below it, but the beginning is smudged and Juliet can only read the last digits “6300”. She instantly recognises Romeo’s number, and decides to go see him. When no-one answers the door, she rushes to a phone booth. She dials 2-1-2-5-2-9-... only to realise she does not remember the final four digits.

First Juliet acts on the information that *Romeo’s number ends in -6300*, later she does not. From the classical perspective, this raises a puzzle. Does Juliet *believe* that Romeo’s number ends in -6300 or not? If she does not, then why did she go to Romeo’s house? But if she does, then why did she not act on this information in the phone booth as well?

The inquisitive picture addresses this puzzle by associating beliefs with a more targeted disposition. A belief is always directed at a particular question, and thus comes with a disposition to act on the belief *only when faced with that question*. Juliet intuitively faces different questions at home and in the phone booth: respectively, *Whose number ends on -6300?* and *What are the last four digits of Romeo’s number?* Her response to the note is guided by her answer to the former question, *Romeo’s*. And her inaction in the phone booth shows she lacks the answer to the latter question, *-6300*. The classical picture gets in trouble by conflating these two truth-conditionally equivalent beliefs. The inquisitive picture distinguishes them, and associates them with different cognitive capacities.

Question-directed belief contents also seem to help with the problem of deduction. This approach to the problem was first suggested by Robert Stalnaker (1991), although he quickly dismissed it. More recently, Seth Yalcin and others have pursued the idea (Yalcin 2008, 2011, 2018, Koralus and Mascarenhas 2013, Pérez Carballo 2016, Yablo 2017). Building on this work, the present paper sets its sights a bit higher, asking not only how deduction is *possible*, but also how it can be *useful*. What is the practical point of gathering beliefs whose informational content we already possess? That is, how does it lead to better choices? This is the *practical* problem of deduction.

This practical problem of deduction is deeply connected to the classical, Ramseyan belief-action principle. That is because the latter gives rise to the classical view of belief states, according to which all our various beliefs cohere into a consistent and deductively closed worldview. To be more precise, a natural articulation of the classical belief-action principle entails that we behave *as if* our beliefs formed such a state. But if that were so, then deductive inquiry would have no practical use. Deductive inferences would make no difference to our choices, since we would already act on the deductive consequences of our beliefs. As shown below, this holds true *even if* we combine the classical belief-action principle with a hyperintensional individuation of belief states at the mental level. So without a more sensitive belief-action principle, there is no practical difference between informationally equivalent doxastic states, and hence no practical point to deductive inquiry.

In particular, this means that the question-sensitive accounts of belief content that Yalcin and others propose do not by themselves address the practical problem of deduction. To do that, we need a story about how question-sensitivity manifests in action. In fact, the difficulty of supplying a general story

of this kind is precisely what prompted Stalnaker to reject the question-directed model of belief:

“[It is not] clear how to generalize [this model] to an account of knowledge and belief in terms of capacities and dispositions to use information (or misinformation) to guide not just one’s question-answering behaviour, but one’s rational actions generally. For we want an account of knowledge and belief, not just for expert systems and people who staff information booths, but for all kinds of agents.” (Stalnaker 1991, 253-4)

This paper responds directly to Stalnaker’s challenge, by building an account of questions in action with all the generality he demands. The result is a systematic account of the role of questions and question-directed beliefs in decision making — to my knowledge, the first of its kind.

Just like its classical cousin, the inquisitive belief-action principle also gives rise to a view of belief states. But as we shall see in Part B, this view is less artificial and idealised than the classical view. In particular, inquisitive agents do not necessarily believe every deductive entailment of what they believe, but they do believe every *part* of what they believe (in the sense of Gemes 1994). Like classical belief states, inquisitive belief states are non-fragmented and holistic. But beliefs are united into a loosely knit “web of questions” rather than a monolithic worldview. Thus the inquisitive picture makes room for agents who, like us, fail to believe some logical consequences of their beliefs, and who may have some inconsistent beliefs. It yields systematic, univocal predictions about the choices of such agents, and thereby lets us assess the practical use of deduction by simply comparing agents’ behavioural dispositions before and after a deductive inference.

Part A. Facing Questions and Having Answers

Choices raise questions, and our beliefs are our answers to those questions — or so I claim. In this part of the paper I develop this idea by examining the initial motivation for it (§I), articulating precise notions of *questions* and *answers* (§II), and saying what it takes for a decision situation to *raise* a question (§III), and for an action to be *guided by* an answer (§IV). This then gives us the formally precise articulations of the classical and inquisitive belief-action principles that will form the basis for our discussion of the problem of deduction in Part B.

I. Two Belief-Action Principles

Traditionally, philosophers, psychologists, decision theorists and economists have characterised the link between belief and action in something like the following way (e.g. Ramsey 1927, p. 159):

Classical Belief-Action Principle: A belief that *p* manifests itself in behaviour as a general disposition to act on *p*.

The ROMEO RECALL case from the introduction raised a difficulty for this principle. In a way, Juliet has the information that Romeo’s number ends in -6300, in that she can recognise the digits. But she has no general disposition to act on this information, since she cannot *recall* those same digits. Cognitive asymmetries of this kind are extremely common. For examples, consult the literature on memory retrieval or the nearest crossword. Here is another case based on Elga and Rayo 2021a:

TRIVIAL TROUBLE: Travelling abroad, Tom is accosted by a fearful sphinx. “Don’t worry,” she says, pinning him gently but firmly to the ground, “I’ve eaten. But I will make you rich if you solve this riddle: name me an English word that ends in the letters -MT.” Tom

racks his brain, but in the end admits defeat. Leaving, the sphinx remarks: “You’re not the brightest bulb, are you now? I could have done this in my sleep!” That night, Tom writes a letter home: “I never dreamt that sphinxes were real!” Then he reads it back incredulously: “D-R-E-A-M-T”.

Tom’s actions in this example intelligible, predictable, and indeed entirely unsurprising. Crossword mavens aside, we *expect* that adult English speakers are able to spell, but may still have difficulty recognising a word based on just a few letters.² Yet we run into trouble when trying to analyse Tom’s behaviour classically. The issue this time is whether or not Tom has the information that “*dreamt*” is spelled *D-R-E-A-M-T* (or, if that seems tautological, the information that */drɛmt/* is spelled *D-R-E-A-M-T*, where */drɛmt/* is the word “*dreamt*,” individuated phonetically). To explain Tom’s response to the sphinx classically, we must say Tom lacks this information, since he fails to act on it. But if we say that, then how can we account for his ability to spell the word correctly in his letter? How can Tom be acting on this information at that later time, unless he had it all along?

Since the classical belief-action principle is widely endorsed, these puzzle cases challenge a wide range of views of belief and agency. In particular, classical decision theorists (of the descriptive kind) say the *extent* to which an agent is disposed to act on a proposition matches their *degree* of belief or credence (Ramsey 1926, Savage 1972, Friston 2013). To account for Tom’s spelling habits, they must attribute to him a high degree of confidence that “*dreamt*” ends in *-MT*. But that predicts, falsely, that Tom can answer the sphinx. Or take Bratman’s theory of agency (1984). On this view, beliefs steer our actions by constraining the formation of plans and intentions — an agent’s plans must be consistent with their beliefs. But why does Juliet’s belief that *Romeo’s number ends in -6300* constrain the plan she makes at home, but not the intentions she forms in the phone booth? Analogous questions arise for any account of belief that endorses some version of the classical principle, including those of Hempel 1962, Dennett 1971, Fodor 1978, Stalnaker 1984, and Schwitzgebel 2002.

I am not saying all these theorists are *wrong*. But I do think that the classical belief-action principle they are working with is a pretty coarse approximation. Moreover, I think I have a refinement on offer that casts some light on its limitations. From the classical perspective, a purely intentional, belief-based explanation for the contrast between Tom’s sphinx-answering and letter-writing behaviour looks to be out of reach. But such an explanation can be had once we note that Tom faces different questions on these two occasions. We just need to refine our belief-action principle to exploit this contrast (here, superscripts indicate the question at which the belief content is directed):

Inquisitive Belief-Action Principle: A belief that A^Q manifests itself in behaviour as a disposition to act on A^Q whenever the agent is confronted with the question Q that the belief is an answer to.

Tom knows *What the final two letters of the word /drɛmt/ are*, and acts on this knowledge when writing the letter. But his beliefs are silent on the sphinx’s question, *What English words end on -MT*.

Some defenders of the classical picture may be skeptical of the explanatory demand to which the inquisitive belief-action principle responds. A common response to cases like TRIVIAL TROUBLE runs like this: Tom has the belief that “*dreamt*” ends on *-MT*, and a disposition to act on this information. But

² As it happens, this particular cognitive asymmetry is extremely well-documented, because psychologists often use word completion as a memory test (see for instance Nelson and McEvoy 1984)

that disposition is weak, and subject to all kinds of masking. In the sphinx exchange, the disposition happens not to manifest itself. But that is just how beliefs are: they sometimes manifest and sometimes not. According to Davidson 1976 for instance, our beliefs manifest only rarely and erratically, and there is no systematic story about when it happens. And the fragmented decision theory of Elga and Rayo 2021a, b says a belief manifests only when separately specified *elicitation conditions* happen to be met, where these conditions are in principle independent of the belief's content.

In my view, such radical weakenings of the classical belief-action link risk throwing out the baby with the bathwater. If beliefs were really that capricious, there would be no telling whether or not they were going to manifest on a given occasion. This belies the fact that we can, and do, predict and explain people's actions in terms of the contents of their beliefs (for details on this point, see Norby 2014 and Hoek 2019, §3.5). Suppose you ask why Mary eats beans every day, and I answer that Mary thinks that eating beans every day will keep her healthy. Intuitively, that makes for a perfectly satisfactory explanation, assuming Mary wants to stay healthy. But if beliefs are only occasionally elicited, shouldn't we expect that Mary will only eat beans on those special occasions? Conversely, we can reliably infer agents' beliefs from their actions. If you see me calmly sipping a cup of tea, you could reasonably infer that I do not believe the tea to be poisoned. But that conclusion would be unwarranted if it were typical for people to fail to act on their beliefs.

Thus our ordinary reasoning about belief and agency presupposes fairly robust, stable doxastic dispositions. The inquisitive picture respects this, and brings out what is constant in the behaviour displayed in our puzzle cases. Throughout *ROMEO RECALL*, Juliet is able to recognise the last digits Romeo's number, and unable to recall them. Likewise, Tom can spell the word "dreamt" throughout *TRIVIAL TROUBLE*. What the sphinx asks him to do is cognitively very different. Properly viewed, then, these cases show us something about the *nature* of doxastic dispositions, not their stability.

II. Beliefs as Answers to Questions

The intelligibility and predictability of Tom and Juliet's responses may be taken as an indication that our ordinary reasoning about belief and behaviour are sensitive to the cognitive distinctions that the inquisitive picture makes. Independent support for that idea comes from the observation that the interpretation of knowledge and belief reports is often sensitive to *focus* (Dretske 1970, Blaauw 2013, Yalcin 2011). For instance, (1) intuitively reports Tom's answer to the question *Which words end on -MT*, while (2) describes his answer to *What are the final letters of the word "dreamt"*:

- (1) Tom thinks/knows that "*dreamt*" ends on -MT
- (2) Tom thinks/knows that "dreamt" ends on -*MT*

Accordingly, (1) strikes us as false in the context of *TRIVIAL TROUBLE*, since it implies that Tom can recognise the word "dreamt" on the basis of its last two letters, while (2) just asserts that Tom can reproduce those letters. There is a similar contrast between Juliet's knowledge that *Romeo has a number ending in -6300*, versus the belief she lacks, that *Romeo has a number ending in -6300*.

To distinguish the beliefs reported in (1) and (2), I propose we take belief contents to be *question-directed propositions*, or *quizpositions* for short: propositions that are jointly individuated by their informational content and the question they answer. Formally, we can capture this concept as follows:

Def. 1. A question Q is a partition of logical space Ω , the set of all possible worlds. When

two worlds w and v share a cell of this partition, we write $w \sim_Q v$. Any set of Q-cells $A \subseteq Q$ is an *answer* to Q.

Def. 2. A **quizposition**, denoted A^Q , is an ordered pair $\langle Q, A \rangle$ whose first member is the question Q that A^Q is said to be *about*, and whose second member is a Q-answer $A \subseteq Q$. The quizposition A^Q is *true* at a world w if and only if $w \in \bigcup A$.

The idea behind this definition of questions is to characterise a question in terms of the information needed to answer it exhaustively (this is pretty standard; see e.g. Groenendijk and Stokhof 1984).

Applying the concept of quizpositions to TRIVIAL TROUBLE, let S stand for the partition question *How do you spell “dreamt”* and let E be the partition question *Which English words end on -MT*:

$w \sim_S v$ iff the word /drɛmt/ has the same spelling in w and v ,
 $w \sim_E v$ iff at w and v , the same English words end on -MT

The partitions S and E are distinct. Each cell $s \in S$ represents a possible spelling of “dreamt”. Since it could be that multiple words end on -MT, the cells $e \in E$ correspond to possible exhaustive *lists* of words ending in -MT. Let M be the set of all S-cells that represent a spelling for “dreamt” ending in -MT, and let W be the set of E-cells that include “dreamt” on the list. Then the quizpositions M^S and W^E have identical truth-conditions, but answer distinct questions.

III. Facing Questions

Decisions confront us with questions. Selecting a wine for the colloquium dinner, you wonder *Which one will make me look knowledgeable?* Deciding which child to reproach, you face the question *Who lit the curtains on fire?* The link between choices and questions pervades our ordinary decision making: we often describe hard choices in terms of *facing* and *confronting* questions.³ Having spotted this link between questions and choices, we can ask what it consists in.

The analysis I propose is pretty simple. Decisions or choices are standardly represented using payoff matrices, like the one below from Jeffrey 1983:

	Chicken	Beef	Herring
White	1	-1	1
Red	0	1	-1
Rosé	0.5	0	-1

This table represents the decision situation of an agent who is choosing what wine to bring to a dinner party. They want to get a good match, but they are not sure what will be for dinner. The column headings of this table represent the *world states* on which the payoffs depend. Note that these form a partition of logical space — that is, a *question* in our sense. More specifically, it is the question *What’s for dinner*, which is exactly the question that this choice intuitively raises for the agent. So in a sense, the questions we face are already hiding in the formalism of classical decision theory: every major formal treatment of decision theory appeals to world state partitions at the fundamental level (Savage 1972, Jeffrey 1983, Lewis 1981, Joyce 1999). It is just that these partitions have not typically been thought of as *questions*, and their role in determining behaviour has been ignored.

³ Comparable idioms exist in at least the following languages: English, German, Dutch, Italian, Spanish, French, Serbian, Turkish, Mandarin Chinese and Shanghainese. With thanks to my informants Vera Flocke, Simona Aimar, Andrés Soria Ruiz, Louis Rouillé, Milica Denić, a student at Bilkent University and Linmin Zhang.

In order to say something informative about what it is for a choice to raise a certain question, I will need to characterise choices or decision problems a bit differently, so as to avoid building in the world state partition as primitive. Here is a way to do this:

Def. 3. An *option* is a real-valued function $\mathbf{a}: \Omega \rightarrow \mathbb{R}$ from possible worlds to utility values. A *decision problem* Δ is a finite set of options.

A decision problem is an abstract representation of a choice. The decision problem Δ faithfully represents a choice just in case there is a correspondence between the actions that the agent is choosing between and the options \mathbf{a} in Δ , such that the value of $\mathbf{a}(w)$ accurately reflects the desirability, to the agent, of the outcome that, at w , would have been obtained had they performed the corresponding action on this occasion. (For simplicity, I presuppose both a causal approach to decision theory and a Stalnakerian semantics for counterfactuals.⁴)

Payoff matrices specify a decision problem thus defined. Each row represents an option, listing the utility values with respect to the worlds in each column. Note that this representation requires that the column partition only groups worlds together where each option takes a constant value. That is, the column partition Q must have the following property:

Def. 4. The choice Δ *raises* the question Q , or Q *addresses* Δ , just in case for every option $\mathbf{a} \in \Delta$, and every cell $q \in Q$, the outcome $\mathbf{a}(w)$ takes on a constant value for all $w \in q$, denoted ' $\mathbf{a}(q)$ '. An agent *faces* the question Q when they make a choice that raises Q .

In other words, a question addresses a choice just in case any complete answer to the question entails what the outcome of each option would be. Multiple questions can address a given decision situation. In particular, if Q addresses Δ , then any question that forms a more fine-grained partition than Q addresses Δ as well. For if an answer to Q already suffices to entail the payoffs of every option, then any cell in the finer partition will definitely suffice. (This observation will become important in Part B.)

Now consider the decision problem that Tom faces when writing down the word "dreamt", Δ_{Letter} :

	/drɛmt/ is spelled D-R-E-A-M-T	/drɛmt/ is spelled D-R-E-A-M-E-D	/drɛmt/ is spelled D-R-E-M-T	...
write D-R-E-A-M-T	1	0	0	...
write D-R-E-A-M-E-D	0	1	0	...
write D-R-E-A-M-P-T	0	0	0	...
...

The options in Δ_{Letter} , listed on the left, are letter combinations you could write down. The aim, let's say, is to spell the word correctly, so there are just two outcomes: success and failure, 1 and 0. Which of those outcomes results from each action depends on what the correct spelling of the word in fact is. Consequently, the column headings form the partition question S , *How is /drɛmt/ spelled?* That is why Tom's spelling of "dreamt" in TRIVIAL TROUBLE was guided by his belief M^S , that *The word /drɛmt/ ends in -MT*: this belief answers the question that Δ_{Letter} raises.

⁴ Assuming the outcomes of the choice form a partition, the Stalnakerian semantics is needed to guarantee that, at every possible world w and for any action A available in the choice C , it is true of exactly one outcome $O_{A,w}$ that "If the agent were to choose A in response to C , outcome $O_{A,w}$ would occur."

But confronting the sphinx, Tom has a different set of options, and the outcome of his choice depends on a different feature of the world. This time, he is choosing between possible replies. Besides confessing ignorance, Tom could venture some answer, like “unkempt”. In the actual world, the only winning reply is “dreamt”, since that is only English word ending in -MT. But in a world with different spellings, “dreamt” would yield failure, and “prompt” success, as shown in the table below. From the column headings we can see that this decision problem — call it Δ_{Sphinx} — raises E, *Which English words end on -MT?* Thus Tom’s choice is guided by his views on E. Hence we can account for Tom’s failure to produce the correct answer in terms of the fact that Tom had no view on E. Tom’s view on S fails to address Δ_{Sphinx} .

	Only /drɛmt/ ends in -MT	Only /ʌn'kɛmt/ ends in -MT	Only /ʌn'kɛmt/ and /prɒmt/ end in -MT	...
reply /drɛmt/	1	0	0	...
reply /prɒmt/	0	0	1	...
reply /ʌn'kɛmt/	0	1	1	...
...
confess ignorance	0	0	0	...

Thus we can explain the difference between Tom’s response to Δ_{Letter} and Δ_{Sphinx} on the basis of an objective contrast between these choice situations: namely a difference in the question these situations raise, which is explicated as a difference in the pattern of counterfactual dependence between the actions available and the outcomes that would result from them.

IV. Acting on Answers

By making quizpositions the objects of belief, we forged a link between beliefs and questions. In the previous section, we linked questions to choices. Chaining these links together, we get a new way of understanding the relation between beliefs and choices. The final puzzle piece is a precise notion of what it takes to be disposed to *act on* a given belief.

Proponents of the classical picture typically gloss acting on a belief as doing what is best given the truth of the belief, or what best promotes one’s desires (e.g. Ramsey 1926, p. 174). We can capture this idea formally by using the notion of dominance:

Def. 5. Suppose **a** and **b** are options, and p is a proposition. Then option **a** (*strictly*) *p-dominates* option **b** just in case $\mathbf{a}(w) > \mathbf{b}(w)$ for all worlds w at which p is true. An option $\mathbf{a} \in \Delta$ is *p-dominant* just in case **a** strictly *p-dominates* every other option in Δ .

We can then say that a general disposition to *act on* p means foregoing p -dominated options in every decision situation that has such options, and hence *performing* the p -dominant option in any situation where there is one. Besides simple decisions, we should take this to include *composite* decision situations that consist of multiple component choices (as explained in §V, this stipulation is needed to ensure correct handling of situations where an agent acts on multiple beliefs jointly).

This yields the following formalisations of the classical and inquisitive belief-action principles:

Classical Belief-Action Principle (formal). A belief that p manifests in action as a disposition to forego p -dominated options in all decision situations.

Inquisitive Belief-Action Principle (formal). A belief that A^Q manifests in action as a disposition to forego A^Q -dominated options in any decision situation that raises Q .

With these precise statements of the classical and inquisitive belief-action principles in place, the analysis of Tom's TRIVIAL TROUBLE can be run entirely within the formalism. Suppose we want to predict Tom's response to Δ_{Letter} on the basis of this formalised classical principle. Then we must attribute a belief p to Tom such that writing D-R-E-A-M-T is the p -dominant course of action in Δ_{Letter} . Inspection of the payoff matrix shows that p must then entail that /drɛmt/ is spelled D-R-E-A-M-T. But if p entails this, the reply "dreamt" in Δ_{Sphinx} beats "I don't know" at every p -world. Applying the classical principle again, we would then get the incorrect prediction that Tom will give the former reply. So as soon as it secures the right prediction for Δ_{Letter} , the classical account slips into the wrong prediction about Δ_{Sphinx} . The inquisitive belief-action principle, on the other hand, lets us explain Tom's action in Δ_{Letter} by attributing to him a belief in the quizposition M^S . But in the inquisitive setting, that attribution does not commit us to any prediction about Tom's response in Δ_{Sphinx} , since that choice does not raise the question S that this belief M^S is an answer to.

Part B. The Web of Questions

We now have precise articulations of the classical and inquisitive belief-action principles. In this second part of the paper, we will see how these two principles give rise to different accounts of the way that our beliefs come together in belief states, where a *belief state* is the totality of all the beliefs that an agent holds at a time. In particular, we will see how the inquisitive picture addresses one of the chief difficulties facing the classical view of belief states: the problem of deduction.

The classical view of belief states is an extreme form of holism. It combines all of an agents' beliefs into a single, global worldview, which Ramsey called "the map by which we steer" (Ramsey 1926, p. 238; Yalcin 2018). The strength of this view is that it captures the way an agent's beliefs unify their actions across domains. But there is a cost: to fit together into a single worldview, classical beliefs must all cohere. Hence the classical view of belief states attributes unbounded, infallible deductive powers to agents, by assuming that their beliefs are deductively closed and perfectly consistent.

This is obviously unrealistic. If we really believed everything entailed by our beliefs, we wouldn't need calculation or deductive reasoning of any kind. New beliefs would only be formed when we acquired new information, and all their ramifications would be instantly known. Calculators would go unsold, mathematicians would be unemployed, Rubik's cubes would be instantly solvable. There is a real puzzle here. Why exactly *does* a Rubik's cube perplex us? After all, the information needed to unscramble it is right in front of us. Then how can it be so difficult? How can reasoning about the cube help us solve it, when it only yields information we already possess?

The ROMEO RECALL and TRIVIAL TROUBLE cases analysed in Part A hint at a way to address this conundrum. The agents in these stories do not lack any relevant information, they need *answers*. In the phone booth, Juliet needs to know *What the final four digits of Romeo's number are*. As it happens, the answer to that question carries information she already has — but that does not detract from its practical usefulness. Likewise, in TRIVIAL TROUBLE, Tom lacks the answer to the sphinx's riddle, even though he has the relevant information in a different form. So apparently, answers can be valuable even when they carry no novel information.

Below I show how to exploit this feature of the inquisitive picture to address the practical problem of deduction. First, §V demonstrates the connection between the classical belief-action principle and the problem of deduction, explaining how this result brings out the practical dimension of the problem. Then §VI-VIII articulate the inquisitive view of belief states, and link it to the inquisitive belief-action principle from Part A. Finally, §IX-XI show with examples how this new view can systematically describe and predict the actions of agents who fall short of the classical ideal, and how it explains what such agents stand to gain, practically speaking, from deductive reasoning.

V. The Practical Problem of Deduction

As noted in Part A, the classical belief-action principle enjoys wide support amongst theorists of belief. The classical view of belief states is considerably less popular. Even its apologists typically concede that a measure of idealisation is involved in the assumption that human beings never have inconsistent beliefs, for instance. What is less widely appreciated, however, is that there is a tight conceptual connection between these two components of the classical picture, so that any criticism of the latter also draws scrutiny to the former.

Let's begin with a statement of the classical view of belief states. As is standard in decision-theoretic contexts, I set aside complications arising with infinities by assuming that the background space Ω of possible worlds is finite. Then the classical view of belief states can be stated as follows:

Def. 6. A classical information state is a set of propositions \mathbf{I} such that:

- i) *Closure under entailment / necessitation*: If $p \in \mathbf{I}$, and if q is true at all possible worlds where p is true, then $q \in \mathbf{I}$.
- ii) *Closure under conjunction*: If $p, q \in \mathbf{I}$, then their conjunction $(p \wedge q) \in \mathbf{I}$.

An information state \mathbf{I} is *accurate at* a possible world w if and only if all propositions $p \in \mathbf{I}$ are true at w ; \mathbf{I} is *consistent* if and only if it is accurate at some world.

Classical Belief States. An agent X 's beliefs form a consistent classical information state \mathbf{B}_X and manifest as a general disposition to forego $\bigwedge \mathbf{B}_X$ -dominated actions.

One can, without loss, represent a classical belief state \mathbf{B} thus defined as the non-empty set of possible worlds where all of the agent's beliefs are true — these are called the agent's *belief worlds*. Using this concept, the classical view of belief states can be tidily summarised thus: agents *believe* whatever is *true* at all their belief worlds, and *do* whatever is *best* at all their belief worlds.

As noted above, the classical assumptions of deductive closure and consistency do not hold true of the belief states of real-world agents. To examine the problem this poses, consider a simple example of deductive failure, from the behavioural economics literature:⁵

MITTEN STATE MURDERS: You ask Mandy, a smart Arizona criminology major, how many murders took place in Michigan last year. She hesitantly guesses "around 150". Then you ask "What about *Detroit*, Michigan?", to which she replies: "I didn't think of Detroit! That city alone had over 200 murders last year. So Michigan's number is much higher."

⁵ Kahneman and Frederick 2002. In this experiment, a group of students from the University of Arizona were asked to estimate the yearly murder rate in Detroit, and another group was asked to estimate the rate in Michigan. The median response for "Detroit" was 200, for "Michigan" 100. (The actual rates are higher.)

Let's attempt to analyse the case from a classical viewpoint. What is the world like according to Mandy? From her confident reply to the follow-up question, it is apparent that Mandy knew that there are over 200 murders a year in Detroit, which is in Michigan. Clearly she did not learn that during this conversation: she had the information all along. But on the classical picture, we are then forced say that Mandy believed from the start that there are over two hundred murders a year in Michigan. For if all of Mandy's belief worlds are worlds where Detroit, in Michigan, has 200+ murders, then there are 200+ murders in Michigan at all of her belief worlds. But on the basis of *that* belief you would predict, incorrectly, that Mandy should give a reply to that effect.

But the difficulty for the classical picture runs deeper. Suppose we weakened or dropped the assumption of closure under entailment, and allowed for the possibility that Mandy believes that *Detroit, Michigan, had 200+ murders last year* but not that *Michigan had 200+ murders*. Then we still could not correctly predict her replies. For the belief that *Detroit, Michigan had 200+ murders* is by itself sufficient to establish that Mandy's reply "around 150" is strictly dominated (for instance by the reply "around 300"). So the classical belief-action principle already makes the wrong prediction about Mandy's behaviour, simply on the strength of her Detroit belief.

This follows a more general pattern. Let p be some arbitrary proposition that Mandy believes, and let q be any proposition that is entailed (necessitated) by p . The disposition classically associated with believing q is this: Mandy avoids q -dominated options whenever they arise (§ IV). Now let \mathbf{b} be any option that is q -dominated by an alternative \mathbf{a} , so that $\mathbf{a}(w) > \mathbf{b}(w)$ at all q -worlds w . Then $\mathbf{a}(w) > \mathbf{b}(w)$ at all p -worlds w : because p entails q , any p -world is also a q -world. And thus Mandy will forego the option \mathbf{b} just on the basis of her belief that p . In general, the classical belief-action principle *by itself* predicts that Mandy behaves in all respects as if she believes any proposition that is entailed by one of her beliefs. At that point, we might just say that she *does* believe those propositions. Closure under single-premise entailment is not an optional posit that is arbitrarily bolted onto the classical view. It is a natural consequence of the belief-action principle at the heart of the classical picture.

The same turns out to be true for the other classical coherence constraints. Since we took the classical belief-action principle to cover composite as well as simple choices, it entails that agents behave as if they believed any conjunction of their beliefs,⁶ and also that having inconsistent beliefs is impossible

⁶ *Proof.* This follows from a Dutch book argument. Suppose an agent is *not* disposed to avoid pq -dominated options. That is, suppose they at least sometimes choose an option \mathbf{b} , even though there is an alternative \mathbf{a} such that $\mathbf{b}(w) < \mathbf{a}(w)$ at all w where p and q are both true. Now suppose that, having chosen \mathbf{b} over \mathbf{a} , this agent is offered a bet. If they refuse, they are guaranteed 0 utility (call that option \mathbf{o}). But if they take the bet (option \mathbf{t}), they win a small prize ε if p is true, and pay a cost C if p is false (pick ε smaller than the minimal difference between \mathbf{a} and \mathbf{b} at pq -worlds, and let C exceed the maximal excess of \mathbf{b} over \mathbf{a} at a q -world). If our agent avoids p -dominated options, they will take the bet, as \mathbf{t} strictly p -dominates \mathbf{o} . But the composite choice of \mathbf{b} and \mathbf{t} is strictly q -dominated by the composite choice of \mathbf{a} and \mathbf{o} :

$$\mathbf{b}(w) + \mathbf{t}(w) = \begin{cases} \mathbf{b}(w) + \varepsilon < \mathbf{a}(w) = \mathbf{a}(w) + \mathbf{o}(w) & \text{for all } w \text{ where } q \text{ is true and } p \text{ is too} \\ \mathbf{b}(w) - C < \mathbf{a}(w) = \mathbf{a}(w) + \mathbf{o}(w) & \text{for all } w \text{ where } q \text{ is true and } p \text{ is not} \end{cases}$$

So in taking the bet, our agent performs a q -dominated option $\mathbf{b} + \mathbf{t}$. But leaving the bet is p -dominated. No matter what they do, having chosen the pq -dominated option \mathbf{b} , the agent can no longer avoid both p -dominated and q -dominated options. Contrapositively, if an agent *does* avoid both p - and q -dominated options, it follows that they avoid pq -dominated options too. ■

for a classical agent.⁷ Putting it all together, the classical belief-action principle entails that agents behave *as if* their beliefs formed a consistent, classical information state.

Granted, that result still leaves a little bit of light between the classical belief-action principle and the classical view of belief states. While the former entails that agents behave *as if* their beliefs formed a consistent, classical information state, the latter says agents' beliefs actually *do* form such a state. One can bridge this gap with a principle like the following:

(Quacks Like a) Duck Principle. If an agent X has the behavioural dispositions that are associated with a belief with a certain content, and moreover X has those dispositions in virtue of their beliefs, then X actually does have a belief with that content.

If an agent looks like they believe p , swims like they believe p , quacks like they believe p , and does all those things in virtue of their beliefs, then the agent probably believes p . Together with the classical belief-action principle, the Duck Principle entails the classical view of belief states.

Thus critics of classical belief states must reject either the classical belief-action principle as formulated in §IV, or reject the Duck Principle. Now if our target is the practical problem of deduction, the latter option is moot. Rejecting the Duck Principle does not shield us from any of the behavioural consequences of the classical view of belief states, or from the conclusion that inconsistent beliefs are impossible — these follow from the classical belief-action principle by itself. Mandy's actions in *MITTEN STATE MURDERS*, and the behaviour of a person attempting to unscramble a Rubik's cube, show that people's beliefs are not classically coherent. But they also show, more directly, that agents do not act *as if* they had classically coherent belief states. Ordinary human agents do not look classical, swim classical or quack classical. That observation conflicts with the classical belief-action principle directly, whether you like the Duck Principle or not.

The problem of deduction is typically discussed in the context of doxastic logic, where it turns into a search for suitably weakened closure principles. As we now see, the problem looks quite different in the practical context. We know at the outset that replacing the classical closure conditions with weaker ones, or even nixing them altogether, will not by itself make any difference to the range of behaviours we can predict and explain. Weakening the closure conditions lets you say that, after deduction, the agent *believes* new things about the cube. But it does not explain how deduction enables agents to *do* new things. The only way we can make progress on this practical problem of deduction is at the roots. We are forced to re-examine the link between belief and action.

As fortune would have it, that is just what we did in Part A of this paper. Over the next two sections, I explain how the inquisitive belief-action principle articulated there gives rise to weaker doxastic closure conditions, and thus yields a principled new view of belief states. The exposition runs a little backwards, starting at the destination. §VI covers the technical preliminaries necessary to state inquisitive view of belief states. Once we have a statement, §VII then shows how this view can be derived on the basis of the inquisitive belief-action principle.

⁷ *Proof.* Suppose, for contradiction, that a classical agent had inconsistent beliefs p_1, p_2, \dots, p_n . By the proof above, it would follow that this agent avoids \perp -dominated options, where \perp is the necessary falsehood $p_1 \wedge \dots \wedge p_n$. But it is not possible to avoid \perp -dominated options. In a binary choice $\{\mathbf{a}, \mathbf{b}\}$, \perp entails both that \mathbf{a} strictly dominates \mathbf{b} , and also that \mathbf{b} strictly dominates \mathbf{a} . So any choice the agent makes is a \perp -dominated option. ■

VI. Quizpositional Mereology

Frege said that some entailments are contained in a proposition the way beams are contained in a house, while others are more like plants contained in their seeds (Frege 1884, §88). The step from *Austin lives on 21 Broad Street* to *Austin lives on Broad Street* is a typical beam-in-house entailment. But the step from *All firefighters are tall* to *No short Canadians are firefighters* is a plant-in-seed entailment. In line with Gemes 1994, Yablo 2017 and Fine 2017, let's call the beam-in-house entailments of a proposition its *parts*. According to the inquisitive view of belief states, we do not believe all deductive consequences of our beliefs, but we do believe every *part* of what we believe (Yablo 2017 and Hawke 2016 defend an analogous view for knowledge).

To see what this means, we must define quizpositional *parthood*, which in turn requires us to say more about questions. First, the conjunction of two partition questions is their coarsest common refinement:

Def. 7. The **conjunction** of two questions Q and R , written QR , is the question

$$QR = \{ (q \cap r) : q \in Q \text{ and } r \in R \} \setminus \{\emptyset\}$$

Equivalently, QR is the partition such that $w \sim_{QR} v$ if and only if $w \sim_Q v$ and $w \sim_R v$.

For instance, a complete answers to the conjunctive question *How many daughters did Russ have and how many sons?* combines a complete answer to each conjunct. A conjunctive question partition is a refinement of each of its conjuncts (in fact, it is always the coarsest partition that has that property). In general, when one question is a refinement of another, we say the former contains the latter:

Def. 8. One question Q **contains** (or is **at least as big as**, or **entails**) another question R if and only if every R -cell is a union of Q -cells. R is **part of** Q if and only if Q contains R .

Equivalently, R is part of Q just in case $w \sim_R v$ whenever $w \sim_Q v$.

For example, *What month is it?* is part of *What date is it?* Note that Q contains R if and only if $QR = Q$. Any conjunction of parts of Q is itself a part of Q . Hence the common parts of two questions are closed under conjunction, which means there is always a greatest common part:

Def. 9. The **overlap** (or **meet**) of two questions Q and R is the biggest question that is both part of Q and part of R . Two questions **overlap** if and only if their overlap is not equal to the empty question $\{\top\}$.

Quizposition conjunction is naturally defined in terms of question conjunction:

Def. 10. The **conjunction** of a Q -answer A and an R -answer B , written AB , is the QR -answer $\{(a \cap b) : a \in A \text{ and } b \in B\} \setminus \{\emptyset\}$. The conjunction of the quizpositions A^Q and B^R , written AB^{QR} or $A^Q \wedge B^R$, is the quizposition $\langle QR, AB \rangle$.

A quizpositional conjunction makes just enough distinctions between possible worlds to make every distinctions made by its conjuncts, and rules out just enough possibilities to rule out every possibility ruled out by its conjuncts. We define quizpositional **negation** and **disjunction** thus: $\neg A^Q := \langle Q, Q \setminus A \rangle$, and $A^Q \vee B^R := \neg(\neg A^Q \wedge \neg B^R)$.

The notion of quizpositional parthood is again a generalisation of the relationship quizpositional conjuncts bear to their conjunction. Thus one quizposition is part of another if it makes no more distinctions *and* rules out no more possibilities:

Def. 11. A quizposition A^Q **contains** a quizposition B^R if and only if Q contains R and

A^Q entails B^R (that is, $\bigcup A \subseteq \bigcup B$); alternatively, we can say B^R is *part of* A^Q . If R is any part of Q , the *maximal R-part* of A^Q , written A^Q/R , is the part of A^Q about R that contains all other parts of A^Q about R .

As in the case of questions, one quizposition contains another just in case their conjunction is equal to the whole. That is to say, A^Q contains B^R iff $AB^{QR} = A^Q$.




FIGURE 1: QUIZPOSITIONAL PARTS

Figure 1 illustrates the definition of parthood visually. Each of these three squares represents a quizposition. The black lines partitioning each square represent its question component. The colouring represents the truth-conditions: bright green for cells where the quizposition is true, and dark red for cells where it is false. The top quizposition contains the two bottom quizpositions: the parts answer more fine-grained questions than the whole, and are true at fewer worlds. (In fact, these are both maximal parts, since if they ruled out any more cells, they would no longer be entailed.)

Doxastic closure under parthood codifies the intuitive idea that in answering a big question, one thereby also answers any parts of that question (see also Hoek *fc.*). For instance, if you believe that *Jane's address is 23 Mountain Drive*, you believe that *Jane's street is Mountain Drive*. Conversely, your views on a big question incorporate your answers to its parts. If you believe that *It is the 20th* in answer to *What day of the month is it*, then you cannot simultaneously be unsure whether *It is the 20th or the 21st of April* in answer to *What date is it*. We can capture the latter idea with a restricted conjunctive closure principle. If you believe both A^Q and B^R , and R is part of Q , then B must also be part of your view of Q . So you also believe the conjunction AB^Q .

To define inquisitive information states, we replace *Closure under entailment* in the classical definition with *Closure under parthood*, and replace *Closure under conjunction* with this restricted version:

Def. 12. An *inquisitive information state* is a set of is a set of quizpositions I subject to the following closure conditions:

- i) *Closure under parthood*: If $A^Q \in I$ and A^Q contains B^R , then $B^R \in I$.
- ii) *Partial closure under conjunction*: If $A^Q, B^R \in I$ and Q contains R , then $AB^Q \in I$.

An inquisitive information state I is *consistent* if and only if there is a possible world at which all quizpositions in I are true. I is *coherent* just in case it contains no necessary falsehoods (i.e. no quizposition of the form $\perp^Q = \langle Q, \emptyset \rangle$). The *domain* of I , denoted \mathcal{D}_I , is the set of all questions about which I contains at least one quizposition. For any $Q \in \mathcal{D}_I$, I 's *view* on Q , denoted $I(Q)$, is the strongest quizposition V^Q in I that is about Q .

Note that partial conjunctive closure ensures that, for any question $Q \in \mathcal{D}_I$, I contains the conjunction

\forall^Q of all quizpositions about Q in \mathbf{I} — this guarantees that $\mathbf{I}(Q)$ must be well-defined for any $Q \in \mathcal{D}_I$. Now we can state the inquisitive view of belief states:

Inquisitive Belief States. An agent X 's beliefs form a coherent inquisitive information state \mathbf{B}_X , and manifest themselves in a disposition to forego $\mathbf{B}_X(Q)$ -dominated options when confronted with any question $Q \in \mathcal{D}_{\mathbf{B}_X}$.

The next section shows how this claim follows naturally from the inquisitive belief-action principle, just like the classical coherence constraints follow from the classical belief-action principle.

VII. Inquisitive Holism

Someone who can tell you the date can always tell you the month. And anyone who can tell you my phone number can tell you its second digit. This is no coincidence. Naming the month is *part* of giving the date, and telling you the second digit is *part* of telling you my phone number. That is why, in each case, the latter ability entails the former. In general, tasks are often composed of smaller subtasks, so that an ability to perform the larger task requires an ability to perform the subtasks.

The mereology of beliefs developed in § VI parallels this intuitive mereology of tasks. Suppose an agent performs a complex task consisting of several subtasks. Suppose also that the agent has beliefs about the question that the larger task confronts them with. Then the inquisitive belief-action principle says that the agent's response to the larger task will be guided by those beliefs. But it *also* says that their response to each subtask is guided by their beliefs about the smaller question that this subtask confronts them with. On pain of contradiction, the guidance from the agent's view on the big question must therefore be in harmony with the guidance coming from their views about its component questions. Assuming the Duck Principle from § V, this turns out to imply that an agent's belief state must satisfy all three of the inquisitive coherence conditions formulated in § VI: it must be closed under parthood, "partially" closed under conjunction, and coherent.

The key to establishing this result is an observation from § III: if some question R is big enough to address a certain decision problem Δ , then it follows from definition 4 that any question that is more fine-grained than R also addresses Δ . In other words:

If the question R is part of the question Q , then Q addresses every decision problem that R addresses.

From this, we can infer that if Q contains R , and an agent has views on both Q and R , then all the choices that are guided by the agent's views on R are also guided by their view on Q .

Let's first think about this informally. Let Q be the question *What are the two biggest cities in Brazil, in order?* Suppose some agent, Abby, has a view on Q , namely that **Rio and São Paulo are the two biggest cities in Brazil**. But she takes no stance on which is bigger. How does this constrain her views about smaller questions? In particular, what does it tell us about Abby's view on R , *what is the biggest city in Brazil?* In figure 2 below, A^Q represents Abby's view about Q . The green cells are the two possibilities she considers live for practical deliberation, namely *Rio is the biggest and São Paulo the second-biggest* and *São Paulo is the biggest and Rio the second-biggest*. The quizpositions B^R , C^R , D^R and E^R represented possible views on R . Our task is to investigate which of these views could in principle be Abby's.

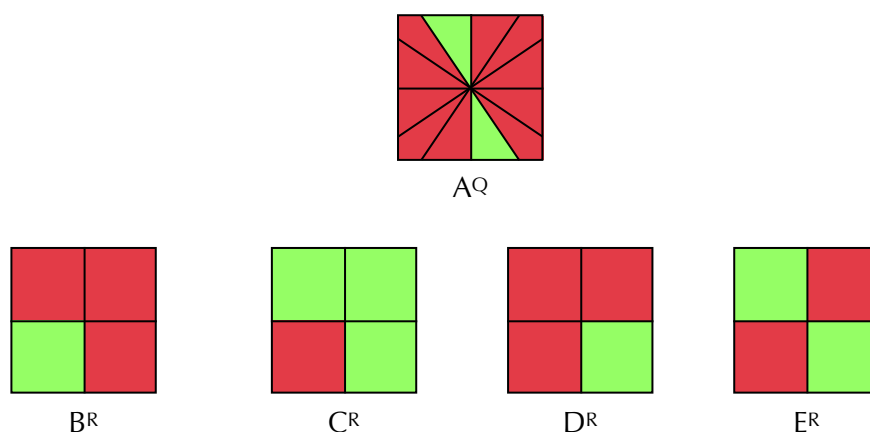


FIGURE 2: WHICH VIEWS OF THE PART FIT THE VIEW OF THE WHOLE?

Firstly, Abby's view on R must be consistent with A^Q . It cannot be, say, B^R , that *the biggest city in Brazil is Salvador*. That's because the behavioural dispositions associated with A^Q and B^R are inconsistent. For instance, when asked whether or not Salvador is the biggest city Brazil, someone with the view A^Q would be disposed to answer "No", and someone with the view B^R would answer "Yes". Abby cannot be disposed to do both. Similarly, we can argue that Abby's view on R must rule out every possibility that A^Q rules out. For she cannot simultaneously act in a way that hedges on a given R-possibility and also in a way that conclusively rules it out. So her view on R cannot be C^R , *the biggest city in Brazil is not Salvador*. For the same reason, her view on R must treat as live every possibility than A^Q treats as live — so it can't be D^R , that *the biggest city in Brazil is São Paulo*. This leaves only one possibility, namely the view that rules out all and only those R-cells that A^Q rules out. This is the view $E^R = A^Q/R$, that *the biggest city in Brazil is either Rio or São Paulo*.

The fact that Abby's view on R rules out every R-possibility that A^Q rules out corresponds to *Closure under Parthood*. The fact that it rules out *only* R-possibilities that A^Q rules out corresponds to *Partial Conjunctive Closure*. Hence these two closure conditions together imply that an inquisitive agent's view about a big question fully determines their view about every part of that question, in the following way: if R is part of Q and the agent's view on Q is V^Q , then their view on R must be V^Q/R , the view that rules out all and only those R-possibilities that A^Q rules out (definition 11).

We can formally prove the link between the inquisitive believe-action principle of § IV and the inquisitive coherence conditions of § VI. Start with *Closure under Parthood*. Let B^R be any part of A^Q . Suppose an agent X believes A^Q , and hence avoids A^Q -dominated actions in any choice that raises Q. We need to show that X also avoids B^R -dominated options when facing R. So suppose X makes a choice Δ that raises R, and suppose there is an option **a** in Δ that B^R -dominates some alternative **b** in Δ . By definition 4, Q also addresses Δ because Q contains R. Furthermore, **a** A^Q -dominates **b**, since $\bigcup A \subseteq \bigcup B$. So since X avoids A^Q -dominated actions when faced with Q, X foregoes option **b**. Thus X has the disposition associated with believing B^R . By the Duck Principle, X does believe B^R .

Next up is *Partial Closure under Conjunction*. Let R be some part of the question Q. In much the same way as in § V, we can show that no agent who fails to avoid A^Q -dominated options when faced with Q can succeed in avoiding both A^Q -dominated options when faced with Q and also B^R -dominated

options faced with R .⁸ Now suppose some agent X believes both A^Q and B^R , and hence does succeed on both fronts. Then X does avoid AB^Q -dominated options faced with Q . So by the Duck Principle, X believes the conjunction AB^Q . Finally, *Coherence* is the condition that X does not believe a contradictory quizposition \perp^Q . If X believed \perp^Q , it would follow X avoids \perp -dominated options when faced with Q , which is impossible for the same reason as before: *all* options are \perp -dominated.

VIII. Doxastic Daisy Chains

The inquisitive view of belief states gives us a way of understanding deductive inference in terms of building connections between our beliefs. To see how that works, we should first attend to the ways in which inquisitive beliefs are linked, according to this view.

Let's begin with an observation made in the previous section: if a question S is part of a bigger question Q , then an agent's view on Q rules out the same S -possibilities as their view on S does. From this, it follows that, when an agent has views on two questions Q and R that share a part S , those views rule out the same S -possibilities as each other:

Overlapping Views. If I is an inquisitive information state, and two questions $Q, R \in \mathcal{D}_I$ have a common part S , then $I(Q)/S = I(R)/S = I(S)$

That is, whenever two inquisitive views concern overlapping questions, they agree on the shared part. Say you have a view on *What the capitals of Europe are* and also a view on *What the capitals of Asia are*. Then *Overlapping Views* says that those two views agree on the capitals of Turkey and Russia. That is, it cannot be that you believe the capital of Turkey is Istanbul with respect to the former question, and Ankara with respect to the latter. Consequently, some changes in view about *the capitals of Europe* also affect your view on *the capitals of Asia*: namely anything in the overlap. So you can think of the agent's view on *the capital of Turkey* is as quite literally being a shared part between these two larger views: changing the common part affects both wholes to which it belongs.




FIGURE 3: A DAISY CHAIN OF INTERLOCKING VIEWS

⁸ The argument is analogous to that in footnote 6. Suppose our agent X fails to avoid AB^Q -dominated options faced with Q . Then X will sometimes pick \mathbf{b} over the alternative \mathbf{a} in a Q -raising choice Δ where $\mathbf{b}(q) < \mathbf{a}(q)$ for all $q \in AB$. On this occasion, we offer X a bet \mathbf{t} that yields some small utility ε if B^R is true and great disutility $-C$ if B^R is false. Note R addresses $\{\mathbf{o}, \mathbf{t}\}$, and \mathbf{t} B^R -dominates \mathbf{o} . But to take the bet would be to choose $\mathbf{b} + \mathbf{t}$ which is A^Q -dominated by $\mathbf{a} + \mathbf{o}$. Moreover, the composite decision problem consisting of Δ and $\{\mathbf{o}, \mathbf{t}\}$ is addressed by QR , and $QR = Q$ because R is part of Q . So if they take the bet, X chooses an A^Q -dominated option faced with Q , while to leave the bet is to choose a B^R -dominated option when faced with R .

But even when two views concern non-overlapping questions, they need not be independent. For they might overlap with a third view without overlapping one another, as illustrated in Figure 3. Each coarse-grained view on the lower tier represents the overlap between the two bigger views directly above it. The top left and top right views do not overlap. But since they both overlap with the view in the centre, they are not independent from one another either. In general, views on disjoint questions may be linked by one or more daisy chains of intermediate views, where each link in the chain overlaps its neighbours. A change in view at one end can percolate throughout the daisy chain.

To substantiate that dynamic claim, let me define a simple notion of belief update:

Def. 13. The **update** of an inquisitive information state \mathbf{I} by a quizposition A^Q , written $\mathbf{I} + A^Q$, is the smallest inquisitive information state containing $\mathbf{I} \cup \{A^Q\}$ as a subset.

Inquisitive updates model the transition that occurs when an agent acquires a new question-directed belief, while retaining all of their old beliefs.⁹ (Note that a prior belief state \mathbf{B} can be updated with A^Q only if $\mathbf{B} + A^Q$ is still *coherent* in the sense of definition 12. As in the classical case, the acquisition of certain beliefs would require belief revision, and not a simple update.)

Updating a belief state with a quizposition A^Q affects the agent's views on Q and any questions that contain Q . But since inquisitive beliefs are linked together, an update could in principle affect an agent's view on any question, as long as it is linked to Q through a chain of intermediate beliefs. Suppose you know that *the meeting is at three o'clock*. You look at your watch and see that *it's two o'clock*. Would you instantly realise that *you have an hour until your meeting*? On the inquisitive picture, it depends. If the three questions *What time it is*, *What time the meeting is* and *How long it is until the meeting* are appropriately linked, then updating your view on the first question directly affects your view about the third, with no additional reasoning required. If not, the entailment may well escape your attention.




FIGURE 4: AN INQUISITIVE UPDATE ON A DAISY CHAIN

⁹ The idea that beliefs arrive in response to a particular question is independently motivated by recent work in epistemology and psychology, incl. Friedman 2013, 2017, Koralus 2014, Carruthers 2018, Drucker 2020. On the view emerging from this literature, beliefs are the result of an *inquiry* into a particular question. They address whatever question the inquirer was wondering about or paying attention to.

To see how that works in the abstract, consider Figure 4 above. The top tier displays the daisy chain of views also shown in Figure 3. The bottom tier shows how a single update can force a change of view all along the chain. The belief state is updated with a quizposition D^Q that rules out a further cell of the agent's view on the leftmost question Q . By *Overlapping Views*, this change in view also rules out an R -cell, since Q and R overlap. Similarly, the change in view on R comes with a change in view on S . Hence the update directly affects the agent's view on S , even though S does not overlap with Q .

So while an inquisitive agent's views on different questions are less closely tied than classical beliefs, they are clearly not compartmentalised fragments, as in Yalcin 2018, 2021. An ordinary human being has a vast number of views on all sorts of interrelated questions, and likely none of those questions are disconnected from the others. Our beliefs form a complex mereological structure: a *web* of belief. A typical belief web may contain tightly knit hubs of thematically connected views, which are better integrated with one another than they are with beliefs outside the hub. But there will typically also be daisy-chains that connect the various islands. Hence there isn't any principled way to identify isolated fragments of belief in the web.

The more questions an agent's belief state has in its domain, the better connected their web will be, and the more coherent their beliefs. If the domain of an inquisitive belief state includes every question, it is closed under entailment and conjunction, and also consistent. Thus classical belief states re-emerge in the inquisitive theory as a theoretical limit case.

IX. Failures of Deductive Closure

Now we are ready to return to the practical problem of deduction. Let us revisit MITTEN STATE MURDERS, in which Mandy initially guesses that Michigan has fewer than 150 murders, though her later remarks show that she knew all along that Detroit alone has over 200. To explain these actions, we look to Mandy's views about the question M , *the number of murders in Michigan last year*, and about D , *the number of murders in Detroit, Michigan last year*:

$$\begin{aligned} w \sim_M v & \text{ iff } \text{the number of murders in Michigan last year are the same at } w \text{ and } v \\ w \sim_D v & \text{ iff } \text{the number of murders in Detroit, Michigan, last year are the same at } w \text{ and } v \end{aligned}$$

From Mandy's initial guess it appears that she has no strong antecedent beliefs about M . Maybe she lacks a view on M altogether, or her views about M are weak. Crucially, it is clear she does *not* believe that *Michigan had over two hundred murders last year* — call this quizposition A^M .

What about Mandy's view on D ? From her remark on Detroit we can conclude that Mandy believes B^D , that *there were over two hundred murders in Detroit last year*. Although B^D entails A^M , Mandy can believe B^D without believing A^M , because these quizpositions are about disjoint questions: D and M do not overlap. Thus the inquisitive account can straightforwardly explain for Mandy's behaviour: at the start of MITTEN STATE MURDERS, Mandy believes B^D but not A^M .

At the end of the story, Mandy performs a deductive inference from B^D to A^M . Consequently, she acquires a new belief A^M (though this belief entails nothing new). Besides being able to describe such deductive belief acquisitions, the inquisitive view also explains what they are good for. If someone asks Mandy about Michigan murders again, she won't make the same mistake, thanks to her newly

acquired belief A^M . Mandy is now disposed to deploy her knowledge about Detroit in a wider range of choice situations. That is what makes this deductive inference practically useful to her.

We can model Mandy's inference as an update with a necessary truth — the quizposition MD^{MD} . The inquisitive closure conditions guarantee that this update leads to the conclusion B^D . For by partial conjunctive closure, Mandy's new belief MD^{MD} and her old belief A^M combine into the view A^{MD} . Hence she will also believe B^D , which is part of A^{MD} . This way of describing the evolution of Mandy's beliefs captures the psychologically familiar fact that Mandy's insight into the link between the questions D and M is liable to last. Suppose Mandy gets additional information on the Detroit murder rate after she has already made the link between A^M to B^D . Plausibly, she will at this point instantly adjust her view on the Michigan murder rate accordingly. Our model predicts this: as long as Mandy maintains a view on DM, her views on D and M remain linked in this way.

In general, deductive reasoning can often be modelled as the acquisition of necessarily true beliefs — updates with quizpositions of the form Q^Q . Such updates do not give the agent new information, but they do forge new connections in the web of belief, which renders an agent's information practically deployable in a wider range of choice situations. (For more about this model of deductive inquiry, and connections to psychology literature, see §V-VII of Hoek *fc.*)

X. Inconsistent Beliefs

Inconsistent beliefs are often the result of closure failures. To illustrate how this can happen, consider an alternative ending to Mandy's story:

MITTEN STATE MURDERS REDUX: After Mandy's guess that there were around 150 murders in Michigan, you reply: "That is right! In fact, with only 120 murders, Michigan has one of the lowest murder rates of any state." Mandy unreflectively takes your word for it, and repeats the statistic to a fellow student later that day.

This time, Mandy acquires the belief C^M , that *Michigan had 120 murders last year*, and ends up acting on it. If she had made the link with Detroit, she may not have accepted C^M so easily. Most likely, if someone mentioned Detroit in this context, Mandy would realise you gave her false information. If so, Mandy retained the belief about Detroit she started with, B^D . So she now has inconsistent beliefs: there is no possible world at which B^D and C^M are both true.




FIGURE 5: PAIRWISE INCONSISTENT BUT COMPATIBLE BELIEFS

While Mandy's beliefs are *inconsistent*, they are not *incoherent* in the sense of definition 12: Mandy does not believe outright contradictions. This is possible because her beliefs are not fully closed under conjunction. Mandy believes *Michigan had 120 murders last year*. She also believes that *Detroit had over 200 murders last year*. But she does not accept their conjunction, *Michigan had 120 murders last year even though Detroit had over 200*. Thus Mandy's belief state at the end of MITTEN STATE MURDERS REDUX

violates all three classical coherence requirements: it is inconsistent, and closed under neither entailment nor conjunction. Still, the inquisitive view produces systematic predictions about how Mandy will act while occupying this doxastic state: she will make choices whose outcome depends on the Michigan murder rate on the basis of her answer *120*, and choices that turn on the Detroit murder rate on the basis of her answer *over 200*.

If Mandy sees the conflict between her beliefs and revises them to restore consistency, the contrariety in her behavioural dispositions will also be eliminated. As in the original MITTEN STATE MURDERS, we can understand this deductive accomplishment as the result of Mandy combining her views A^M and B^D into a unified view about MD. But this time the transition is not a simple update: it involves belief revision as well. Yet the upshot is similar, in that the views in Mandy's belief web becomes better connected as a result, which helps Mandy act in a more cohesive way.

XI. Necessary Truths and Dutch Books

Sections IX and X showed, using examples, how deductive inquiry leads to more cohesive behaviour. We can spell out this improvement in more exact and general terms by invoking Dutch books. In both MITTEN STATE MURDERS and MITTEN STATE MURDERS REDUX, a synchronous Dutch book could be made against Mandy. In both scenarios, Mandy considers a low Michigan murder rate a live possibility, while ruling out a low Detroit murder rate. So Mandy is disposed to buy a bet on the former proposition at sufficiently favourable odds, and to bet against the latter proposition at very unfavourable odds (assuming, as is customary in this context, that she is the betting type). Together, that combination of bets yields a guaranteed loss. In both cases, Dutch books of this kind are no longer possible once Mandy links her beliefs about these questions by means of deductive reasoning.

These are instances of a more general pattern. As detailed above, deductive accomplishments can often be understood in terms of acquiring a belief in a necessary truth — in Mandy's case, the quizposition MD^{MD} . Let us say a *Dutch Q-book* is a Dutch book in which every bet is addressed by the question Q. According to the inquisitive belief-action principle, an agent who believes Q^Q is disposed to avoid strictly dominated sequences of bets in any composite decision situation that confronts the agent with Q — in other words, they avoid Dutch Q-books. So the behavioural manifestation of believing a necessary truth is to avoid a certain kind of Dutch book.

This observation gives us a more systematic understanding of the way that deductive reasoning improves behavioural coherence, on the inquisitive picture. An agent who believes few necessary truths, and whose belief state therefore has a small domain, may be easily Dutch bookable, since their views are relatively disconnected. At the other end of the spectrum is an agent who believes every necessary truth, and thereby avoids Dutch books altogether: this is the classical agent. The rest of us are somewhere in between those extremes. But we can improve our situation by forming views on new questions. By doing so, we promote the cohesiveness of our beliefs and of our choices, and take incremental steps towards the classical ideal.

In closing, it is worth noting that the link to Dutch Books also opens an avenue towards a principled account of partial beliefs or credences as question-directed attitudes, analogous to the account of inquisitive full beliefs given above. It is well-known that, modulo certain assumptions, agents who avoid Dutch books can be represented as having probabilistically coherent credences, with respect to

which they maximise expected utility (de Finetti 1937, Hájek 2005). Likewise, an agent who always avoids Dutch Q-books maximises expected utility whenever they are faced with Q, relative to some uniquely determined set of credences in quizpositions about Q. This observation suggests the prospect of an inquisitive decision theory that captures choice under uncertainty for logically non-omniscient agents (for an implementation of this idea, see Hoek 2019, Ch. 3-4).



Our lives confront us with all manner of questions. Our beliefs are the answers we have to those questions. This understanding of the role of belief differs from the classical view of belief as the possession of undirected information. In this paper I sharpened this contrast, and showed how it gives rise to a principled new way of understanding doxastic states. Rather than being a store of agglomerated information, a belief state is a complex web of views on interconnected questions. This rich model of cognition makes room for deductive inquiry as an activity that is driven by posing new questions, thereby drawing new connections in the web. This activity is fruitful not because it gives us new information, but because it renders the information we have more widely deployable and more integrated, which in turn leads to more cohesive decision making.

References

- Blaauw, Martijn, 2013, "Contrastive Belief." *Contrastivism in Philosophy: New Perspectives*, ed. Martijn Blaauw, 88-100. New York: Routledge.
- Bratman, Michael, 1984, "Two Faces of Intension." *Philosophical Review* 93(3): 375-405.
- Ciardelli, Ivano and Floris Roelofsen, 2015, "Inquisitive dynamic epistemic logic." *Synthese*, 192(6), 1643-87.
- Davidson, Donald, 1976, "Hempel on explaining Action" *Erkenntnis* 10(3): 239-53.
- De Finetti, Bruno, 1937, "La Prévision: Ses Lois Logiques, Ses Sources Subjectives." *Annales de l'Institut Henri Poincaré* 7: 1-6.
- Dennett, Daniel, 1971, "Intentional Systems." *Journal of Philosophy* 68: 87-106.
- Dretske, Fred, 1970, "Epistemic Operators." *The Journal of Philosophy* 67(24): 1007-23.
- Drucker, Daniel, 2020, "The Attitudes We Can Have." *The Philosophical Review* 129(4): 591-642.
- Dummett, Michael, 1975, "The Justification of Deduction." *Proceedings of the British Academy* 59: 201-32.
- Elga, Adam, and Agustín Rayo, 2021a. "Fragmentation and information access." In Borgoni, Kindermann and Onofri (eds.), *The Fragmented Mind*, 37-53. New York Oxford University Press.
- 2021b. "Fragmentation and logical omniscience." *Noûs*, online first.
- Fine, Kit, 2017, "A Theory of Truthmaker Content I, II" *J. of Philosophical Logic* 46(6): 625-702.
- Fodor, Jerry, 1978, "Propositional Attitudes." *The Monist* 61: 501-23.
- Frege, Gottlob, 1884, *Die Grundlagen der Arithmetik*. Breslau: Verlag Wilhelm Koebner.
- Friedman, Jane, 2013, "Question-Directed Attitudes." *Philosophical Perspectives* 27(1): 145-74.
- 2017, "Inquiry and Belief." *Noûs* 53(2): 296-315.
- Friston, Karl, 2013, "Life as we know it." *Journal of the Royal Society Interface* 10, 1-12.
- Gemes, Ken, 1994, "A new theory of content I: Basic content." *J. of Philosophical Logic*, 23(6):595-620.
- Groenendijk, Jeroen and Martin Stokhof, 1984, *Studies on the Semantics of Questions and the Pragmatics of Answers*. PhD Thesis, University of Amsterdam.
- Hawke, Peter, 2016, "Questions, topics, and restricted closure." *Philosophical Studies* 173(10): 2759-84.
- Hájek, Alan, 2005, "Scotching Dutch books?" *Philosophical Perspectives* 19(1): 139-151.
- Hempel, Karl, 1962, "Rational Action." *Proc. of the American Philosophical Association* 35: 5-23.
- Hoek, Daniel, 2019, "The Web of Questions: Inquisitive Decision Theory and the Bounds of Rationality." PhD dissertation, New York University.
- forthcoming, "Minimal Rationality and the Web of Questions." To appear in Kindermann, Van Elswyk and Egan (eds.), *Unstructured Content*, Oxford University Press.
- Jeffrey, Richard, 1983, *The Logic of Decision*. Second Edition. Chicago: University of Chicago Press.

- Joyce, James M., 1999, *The Foundations of Causal Decision Theory*. Cambridge: Cambridge University Press.
- Kahneman, Daniel, and Shane Frederick, 2002, "Representativeness revisited: Attribute substitution in intuitive judgment." In *Heuristics and biases: The psychology of intuitive judgment*, eds. Gilovich, Griffin and Kahneman, 49-81. Cambridge: Cambridge University Press.
- Koralus, Philipp, 2014, "The Erotetic Theory of Attention." *Mind & Language* 29.1: 26-50.
- Koralus, Philipp, and Salvador Mascarenhas, 2013, "The Erotetic Theory of Reasoning: Bridges Between Formal Semantics and the Psychology of Deductive Inference" *Phil. Perspectives* 27(1): 312-65.
- Lewis, David, 1981, "Causal Decision Theory." *Australian Journal of Philosophy* 59(1): 5-30.
- Nelson, D. L. and C. L. McEvoy, 1984, "Word Fragments as Retrieval Cues: Letter Generation or Search through Nonsemantic Memory?" *The American Journal of Psychology* 97(1): 17-36.
- Norby, Aaron, 2014, "Against Fragmentation." *Thought* 3(1): 30-38.
- Pérez Carballo, Alejandro, 2016, "Structuring logical space." *Phil. and Phenomenological Research* 92(2): 460-91.
- Powers, Lawrence H., 1978, "Knowledge by Deduction." *The Philosophical Review*, 87(3): 337-71.
- Ramsey, Frank, 1926, "Truth and Probability." In *The Foundations of Mathematics and other Logical Essays*, 138-54. Ed. R.B. Braithwaite, London: Kegan Paul 1931.
- 1927, "Facts and Propositions." *Proc. of the Aristotelean Society* 7: 153-170.
- Savage, Leonard J., 1972, *The Foundations of Statistics*. Second Revised Edition. New York: Dover Publications.
- Schwitzgebel, Eric, 2002, "A phenomenal, dispositional account of belief." *Noûs* 36: 249-75.
- Stalnaker, Robert, 1984, *Inquiry*. Cambridge, MA: MIT Press.
- 1991, "The Problem of Logical Omniscience I." Repr. in *Context and Content*, 241-54, New York: Oxford University Press.
- Yablo, Stephen, 2017, "Open knowledge and changing the subject." *Philosophical Studies* 174(4): 1047-71.
- Yalcin, Seth, 2008, "Modality and Inquiry." PhD thesis, Massachusetts Institute of Technology.
- 2011, "Nonfactualism About Epistemic Modality." In *Epistemic Modality*, eds. Andy Egan and Brian Weatherson, 295-332. New York: Oxford University Press.
- 2018, "Belief as Question-Sensitive." *Philosophy and Phenomenological Research*, 97(1): 23-47.
- 2021, "Fragmented But Rational." In Kindermann, Borgoni and Onofri (eds.), *The Fragmented Mind*, Oxford University Press.