

Minimal Rationality and the Web of Questions

DANIEL HOEK, PRINCETON UNIVERSITY AND VIRGINIA TECH, DANIELHOEK.COM, CORRECTED VERSION SEPTEMBER 2023.¹

ABSTRACT: This paper proposes a new account of bounded or minimal doxastic rationality (in the sense of Cherniak 1986), based on the notion that beliefs are answers to questions (à la Yalcin 2018). The core idea is that minimally rational beliefs are linked through thematic connections, rather than entailment relations. Consequently, such beliefs are not deductively closed, but they are closed under parthood (where a part is an entailment that answers a smaller question). And instead of avoiding all inconsistency, minimally rational believers only avoid blatant inconsistencies (where some beliefs are blatantly inconsistent when they contradict one another on a particular question). Rather than cohering into a single overall world view, beliefs are more loosely connected in what is best described as a *web of questions*. This view of minimally rational belief naturally gives rise to an account of deductive inquiry on which deductive reasoning is a matter of posing new questions..

Both ordinary and theoretical explanations of human and animal behaviour tend to turn on the assumption that the agent in question has coherent, rational beliefs. What does that assumption amount to, exactly? *Ideal* rationality, the standard typically assumed in doxastic logic and game theory, seems too strong for many purposes. It requires an agent to have perfectly consistent beliefs, and to be logically omniscient in the sense that their beliefs are closed under entailment. But while this can be a useful idealisation, it is often more than we need to assume, and in some cases it is clearly too much. For example, the purchase of a calculator only makes sense if the buyer is *not* ideally rational. And if we are trying to understand the behaviour of someone who is attempting to solve a Rubik's cube, the assumption of ideal rationality is a non-starter: an ideally rational agent would instantly know how to solve the cube, simply by observing its scrambled state. To really achieve ideal rationality would require instantaneous computational powers and an infinite memory. Ordinary, finite creatures like ourselves do not just fall short of that ideal: we do not even come anywhere close.

So there is a theoretical need for a less demanding, more realistic standard of doxastic rationality. This need arises in philosophy and a range of other disciplines where belief-based explanations of behaviour play a role, such as psychology, economics, legal theory, political theory and computer science. This lower standard should be attainable, and for the most part attained, by real-world believers like ourselves, with finite cognitive and computational resources. But it should still be high enough to sustain ordinary and theoretical belief-based explanations of behaviour.

An early advocate for a lowered standard of rationality was Christopher Cherniak (1986). His outline for a theory of what he dubbed *minimal* rationality will be our starting point here. It goes roughly like this. A minimally rational subject may not see every consequence of their beliefs, but they do generally see the *direct* consequences. And while their beliefs may contain inconsistencies, a minimally rational believer does avoid *blatant* inconsistencies. Cherniak's account also has an important dynamic aspect:

¹ To appear in *Unstructured Content*, eds. Egan, Kindermann, Kirk-Giannini and Van Elswyk, Oxford: Oxford University Press 2023. I originally wrote this chapter during a postdoc at Princeton in 2019. There have been some minor changes and improvements since. I would like to thank Anthia Solaki for the conversation that introduced me to Cherniak's work. Thanks also to Justin Bledin, David Chalmers, Cian Dorr, Paul Egré, Kit Fine, Simon Goldstein, Jordan MacKenzie and David Thorstad for their comments on earlier drafts. Thanks to Dane Stocks and Annelisa O'Neal for editorial assistance.

when the need arises, minimally rational agents reliably make straightforward deductive inferences from their beliefs. But the more difficult an inference gets, and the more cognitive resources it requires, the less likely a subject is to perform it.

In developing this view, the challenge is to flesh out the operative notions of a *direct* consequence, a *blatant* inconsistency and the *difficulty* of a deductive inference in ways that steer clear of both triviality and over-idealisation. If we render the notion of minimal rationality too weak, it ceases to have predictive or explanatory value. But if we make it too strong, it could collapse into ideal rationality or something uncomfortably close. There are strong pressures from both sides, which makes this a notoriously difficult balance to strike.

In this chapter, I show that a simple, intuitive solution to this challenge naturally suggests itself once we take on board a conception of belief advocated by Seth Yalcin (2011, 2018) and others: namely the view that the contents of our beliefs are answers to specific questions, and not undirected pieces of information. On the account of minimal rationality I will propose, minimally rational beliefs are linked together by their *thematic* connections rather than their entailment relations. On this view, a minimally rational subject's beliefs are not perfectly integrated like those of an ideally rational subject. But neither are they partitioned into isolated compartments, as in fragmentation theories of belief. Rather, a distinctive and I think cognitively plausible picture of doxastic states arises, according to which all of an agent's views are indirectly connected to one another in a *web of questions*.

The approach I take here is rooted in an unstructured or non-syntactic view of belief content. This is a departure from the norm. At least in the philosophy literature,² nearly all extant accounts of minimal or bounded rationality are built on the assumption that belief contents are imbued with syntactic structure. Given that state of affairs, one could be forgiven for thinking that it is impossible to articulate a plausible notion of minimal rationality unless one embraces the view that belief contents are syntactically structured. This contributes to the impression that unstructured views of belief content only apply to heavily idealised subjects, and are inadequate for more realistic contexts, where notions like minimal rationality become important.

One core aim of this chapter is to help dispel that impression. As I will argue, the theory of minimal rationality proposed here holds its own against the syntax-based competition, and even has some clear advantages. It is more elegant and principled, and makes a sharp division between minimal rationality and irrationality, rather than leaving that boundary arbitrary or vague. I agree that the sets-of-worlds view of belief contents make an inauspicious starting point for a theory of minimal rationality. But in this chapter I hope to show that this critique does not carry over to the new crop of theories of hyperintensional unstructured propositions, such as the views articulated in Yablo 2014, Fine 2016, 2017, Ciardelli, Groenendijk and Roelofsen 2019, and the present volume.

The first half of the chapter concerns the *static* aspects of minimal rationality. Sections I and II describe the more or less familiar difficulties one encounters when trying to combine an intensional, sets-of-

² Computer science has the awareness-based strategy of Fagin and Halpern 1988. See also Sim 1997, Schipper 2015, Franke and de Jager 2011, Egré and Bonnay 2012, Fritz and Lederman 2015. This approach is not essentially syntactic, and has affinities with my proposal below. But it cannot capture the notion of minimal rationality I am aiming for here, because it precludes the possibility of rational inconsistencies (see §2 below).

worlds account of belief content with the notions of a direct consequence or a blatant inconsistency. This takes the form of two puzzles about minimal rationality: one about closure and one about consistency. In light of those difficulties, Section III proposes a hyperintensional notion of belief contents, and outlines a natural account of minimal, static rationality on this basis. Section IV then revisits the two puzzles we started with, explaining how the new account resolves them both.

The second half of the paper is about the dynamic aspect of minimal rationality. Section V introduces a third puzzle about minimal rationality. It argues that the intensional, sets-of-worlds view of belief content faces serious difficulties in making sense of the observation that some deductive inferences are harder to perform than others. Section VI shows how we can understand deductive inquiry as a question-guided endeavour, an idea that flows naturally from the account of minimally rational belief states developed in the first half of the paper. Section VII shows how this model makes sense of the distinction between easy and difficult deductive inferences, identifying three cognitive obstacles that captured by the model: *conceptual* barriers that prevent us from asking certain “novel” questions, *computational* limitations that prevent us from asking certain “big” questions, and *strategic* limitations that prevent us from identifying the right question to ask.

I. A Puzzle About Closure

Suppose I am of the opinion that *it is 8.30pm*. It would be natural to infer from this that apparently, I believe that *it is not 4.30pm*, and that *it's evening*, and that *it is not yet 9 o'clock*. Or suppose Amy thinks *October 31st will be a warm, cloudy Tuesday*. Then presumably she also thinks that *October 31st is a Tuesday*, that *October 31st will be a warm day* and that *it will be cloudy on October 31st*. Thus our ordinary belief attributions bear out the assumption that when you believe something, you also believe some of its entailments (at least if you are rational). Ideal theories of rationality capture this with the requirement that beliefs are closed under single-premise entailment:

Ideal Closure. Whenever a rational agent believes something, they believe all of its logical consequences. That is, if $\phi \models \psi$, then also $B\phi \models B\psi$.

Here “B” is to be read as “X believes that” where X is an arbitrary rational believer.³

But given that our present aim is to capture our ordinary, minimal rationality assumptions, *Ideal Closure* should be rejected. For example, suppose Joe believes that *there are exactly thirteen cartons of eggs in the box, containing a dozen eggs each*. It does not intuitively follow that Joe believes *there are exactly 156 eggs in the box*: Joe can be (minimally) rational without having bothered to make the calculation. A different sort of counterexample to *Ideal Closure* is emphasised in the literature on attention (e.g. Franke and de Jager 2011). Suppose Emma thinks that *the car keys are nowhere in the house*. Does it follow that she believes that *the car keys are not in the second drawer of the little wooden cabinet by the front door*? The answer seems to be “no”: Emma may not have considered that particular possibility. Other counterexamples to *Ideal Closure* involve conceptual limitations, like this one from Stalnaker 1984 (p. 88): King William believed *he could avoid war with France*. But he did not, intuitively,

³ The schema “if $\phi \models \psi$, then also $B\phi \models B\psi$ ” is validated by a range of logics of belief, and also by the standard natural language semantics for belief reports (e.g. Heim and Kratzer 1998, Ch. 12). A version of the Puzzle about Closure arises in both contexts, though the schema means subtly different things in each one: the schematic letters range over different sentences, and different notions of entailment are in play. See appendix.

believe that *he could avoid nuclear war with France*. As an eighteenth-century monarch, William lacked the concept of nuclear war.

So a theory of minimal rationality needs a weakening of *Ideal Closure*:

Minimal Closure. Whenever a rational agent believes something, they believe all of its direct consequences. That is, if ψ is a direct consequence of ϕ , then $B\phi \models B\psi$.

In fleshing out the notion of a “direct” consequence here, we will have to tread carefully: by closing beliefs under direct consequence, we allow in the direct consequences of those direct consequences as well, and their direct consequences. The worry is that we may reach some very indirect consequences in this manner. It is a familiar fact that even very remote consequences can often be reached through a long sequence of simple steps. So before you know, *Minimal Closure* collapses back into *Ideal Closure*.

But such a collapse is not inevitable. The key is to identify a notion of direct consequence that is transitive. As long as any direct consequence of a direct consequence of ϕ is itself already a direct consequence of ϕ , we stay out of trouble. That way, there is no risk of finding any indirect consequences amongst the direct consequences of the direct consequences. With Cherniak, we might think of deductive inferences as carrying a certain cognitive cost. The very easiest inferences do not carry any cost. Instances of the Reiteration rule $\phi \therefore \phi$ should make for uncontroversial examples. This “inference” requires no effort at all: if you believe its premise, ϕ , then ipso facto you also believe its conclusion, which is also ϕ . As long as we are careful to restrict the moniker “direct consequence” to such zero-cost inferences, the feared collapse will not occur. For though a long sequence of low-cost steps can add up to a costly procedure, a long sequence of zero-cost steps still costs nothing (cf. the treatment of System I inferences in Solaki et al. 2019).

For a non-trivial *Minimal Closure* condition, some inference patterns besides Reiteration must be counted as direct. Cherniak conjectures that conjunction eliminations are the easiest kind of inference — that is, inferences of the form $(\phi \wedge \psi) \therefore \phi$ (Cherniak 1986, p. 28). So conjunction eliminations should count as direct, zero-effort inferences if any inferences do.⁴ To appreciate the intuitive pull of this suggestion, just consider a few examples: if Jill believes that *tigers and zebras are striped*, it seems to follow that she believes *tigers are striped*. And if she thinks *John is nasty, brutish and short*, clearly she believes that *John is short*.

The intuition that belief is closed under conjunction elimination is widely attested (see Dretske 1970, Vardi 1986, Jago 2013, Fine 2016, Hawke 2016, Yablo 2017). Like reiterations, conjunction elimination is so straightforward that it seems questionable whether it is properly speaking an inference at all, or really just repetition of what was already said. As Yablo likes to put it: anybody who believes $(\phi \wedge \psi)$ *already* believes ϕ (Yablo 2014, p. 116). Note also that conjunction elimination is transitive in the desired way. By positing that a minimally rational agent believes the conjuncts of their beliefs, we do not risk letting in anything unforeseen: the conjuncts of the conjuncts of a sentence are themselves also conjuncts of the whole sentence.

⁴ I will assume that some non-trivial inferences are zero-effort and automatic. But I should say that Cherniak himself vacillates a little on this point: some remarks clearly imply that bottom-rung, maximally easy inferences are totally automatic, while others suggest they are merely low-effort (see esp. §1.4 and §2.6-7).

So it is plausible that, if there are non-trivial direct inferences at all, conjunction eliminations should be among them. Intuitively, inferences just do not get more immediate than this. If we take that idea on board, then any non-trivial *Minimal Closure* condition should entail the following:

Closure under Conjunction Elimination. When a rational agent believes a conjunction, they also believe its conjuncts. That is, $B(\phi \wedge \psi) \models B\phi$ and $B(\phi \wedge \psi) \models B\psi$.

Summing up, we want a *Minimal Closure* constraint that is intermediate in strength between *Closure under Conjunction Elimination* and *Ideal Closure*. Or in other words, we are aiming for an account of rationality that endorses *Closure under Conjunction Elimination* but rejects *Ideal Closure*.

The Puzzle about Closure arises when trying to combine these desiderata with the traditional unstructured view of belief as a relation between agents and sets of possible worlds, or indeed any view of belief that endorses the following venerable principle of doxastic logic:

Intensionality. If two propositions are logically equivalent, agents believe one just in case they believe the other: if $\phi \models \psi$, then $B\phi \models B\psi$.

The problem is that, given *Intensionality*, *Closure under Conjunction Elimination* is equivalent to *Ideal Closure*. For if ϕ entails ψ , then ϕ is equivalent to $(\phi \wedge \psi)$. So on an intensional view, believing ϕ comes to the same thing as believing $(\phi \wedge \psi)$. But then you can infer ψ from ϕ using conjunction elimination. Given *Intensionality*, all single-premise inferences are instances of conjunction elimination. Thus we cannot reject *Ideal Closure* if we accept both *Intensionality* and *Closure under Conjunction Elimination*. The Puzzle about Closure is the resulting trilemma between rejecting *Intensionality*, rejecting *Closure under Conjunction Elimination* and accepting *Ideal Closure* (this problem is discussed, in one form or other, in Hawthorne 2009, Kripke 2011, Fine 2013, Hawke 2016 and Yablo 2017).

To maintain *Intensionality* in the face of this puzzle, one has to give up the hope of finding a *Minimal Closure* constraint of the kind we just envisioned. For if you embrace the view that conjuncts are direct consequences, you are forced to say that *every* entailment is a direct consequence. That view endorses *Ideal Closure* for minimally rational subjects, which is problematic in view of the abundance of apparent counterexamples. On the other hand, if you think conjuncts are not direct consequences, then it is hard to see what direct inferences could plausibly remain: intuitively speaking, it does not get easier than conjunction elimination. Either way, there is no space for a *Minimal Closure* principle that occupies a comfortable middle ground between *Ideal Closure* and triviality, because *Intensionality* obliterates the distinction between direct and remote consequences. If you take the demand for a notion of minimal rationality seriously, that is a strong reason to reject *Intensionality*.

However, rejecting *Intensionality* is not, by itself, enough to solve the puzzle. For in fairness to the intensional view of belief, it must be said that its “structured” competitors do not directly shed any great light on the distinction between direct and remote consequences either. Syntax-based views of doxastic rationality typically understand rationality as the result of a series of syntactic operations on a belief state. Belief states here are either modelled as a set of sentences or syntactically structured propositions (a *belief box*), or as a set of logically impossible worlds, where those worlds are in turn sets of sentences. Bounded or minimal rationality is the product of applying deductive inference rules of limited difficulty to such a set, involving sentences of limited length, using a limited number of reasoning steps, etcetera. On this picture, ideal rationality is the theoretical limit of the process, where

all inferences have been made. Proposals in this tradition include Cherniak's theory, Eberle 1974, Moore and Hendrix 1979, Konolige 1986, Gaifman 2004, Bjerring and Skipper 2019 and Solaki, Berto and Smets 2019; other syntactic approaches are based on non-classical logic, including Cresswell 1975, 1985, Levesque 1984 and Fagin, Halpern and Vardi 1995.⁵

Views like these can easily capture *Closure under Conjunction Elimination* without collapsing into *Ideal Closure*. The problem is where to take it from there. What other inferences are direct? Is the inference from $(\phi \vee \psi)$ to $(\psi \vee \phi)$ direct? What about the inference from $(\phi \supset \psi) \wedge \phi$ to ψ ? Even if it is clear from the outset that all sorts of boundaries can be drawn here, the view gives us no guidance about which boundary to pick. For this reason, syntactic accounts of minimal closure tend to involve a good amount of arbitrariness or vagueness. Moreover, views in this tradition are maladapted to capture the intuition that, say, *It's evening* is a direct consequence of *It's 8.30pm*. Sentences do not in general seem to bear any simple, uniform syntactic relationship to their direct consequences.

Ideally, we would like a solution to the Puzzle about Closure that yields some insight into minimal rationality, and gives us some principled guidance on how a *Minimal Closure* principle should be formulated. If that is our aim, then simply rejecting *Intensionality* does not cut it: what we need in its place is a view of belief contents that, unlike the intensional and syntactic views, does something to illuminate the relation that propositions bear to their direct consequences.

II. A Puzzle About Consistency

Suppose Joe knows *he has to go to Sarah's birthday party*. And he also knows that *Sarah's birthday is this Wednesday*. But he has not put two and two together yet to form the belief that his Wednesday night is taken. Consequently, when that question arises, he consults his diary. Finding it tells him *he has Wednesday evening free*, he comes to believe that too. If Joe retained his beliefs about Sarah's birthday, his beliefs have now become inconsistent. In this way, deductive limitations inevitably make one vulnerable to inconsistency too.⁶ That is not to say we believe straight-up contradictions. For instance, in Joe's case, it would not be intuitively correct to say he believes that *he has Wednesday evening free even though he has to go to Sarah's birthday party then*.

So while any plausible theory of minimal rationality must countenance the possibility of inconsistent beliefs, that does not mean anything goes. Minimally rational agents cannot be immune to inconsistency, but we do want to rule out *blatant* inconsistencies. As Cherniak argues, minimal rationality must license inferences about the beliefs a subject *lacks* on the basis of the beliefs they *have* (Cherniak 1986, §1.5). So we should reject *Ideal Consistency* but accept *Minimal Consistency*:

⁵ Shouldn't sets of impossible worlds be classified as unstructured contents? The superficial similarity with sets of possible worlds suggests as much, but that appearance is misleading. See for instance Jago 2015 on the similarity between impossible worlds-based and Russellian accounts of propositions. Unlike possible worlds, impossible worlds are typically characterised as collections of sentences, and views like Jago 2013, Solaki et al. 2019 and Bjerring and Skipper 2019 rely heavily on the syntactic character of impossible worlds.

⁶ In fact, this example involves two failures of deductive closure: first, a failure to form the belief that Wednesday night is taken, and later a failure to appreciate that if his Wednesday night were really free, it would follow that Joe did not have to go to Sarah's birthday party on Wednesday.

Ideal Consistency. A rational agent's beliefs are consistent. That is, if $\phi_1, \phi_2, \dots, \phi_n \models \perp$, then $B\phi_1, B\phi_2, \dots, B\phi_{n-1} \models \neg B\phi_n$.

Minimal Consistency. Rational agents do not believe blatant inconsistencies: if $\phi_1, \phi_2 \dots \phi_n$ are blatantly inconsistent, then $B\phi_1, B\phi_2, \dots, B\phi_{n-1} \models \neg B\phi_n$.

Given that a rational agent has a particular belief, *Minimal Closure* allows us to make inferences about what other beliefs they must have. *Minimal Consistency*, on the other hand, allows us to draw conclusions about what beliefs they must lack. To articulate the *Minimal Consistency* principle, the notion of a "blatant" inconsistency must be analysed.

To begin with, the discussion above suggested that outright contradictions cannot be rationally believed, and should be counted as blatant inconsistencies. Then *Minimal Consistency* entails:

Avoidance of Contradictions. Rational agents do not believe contradictions. That is, if $\phi \models \perp$, then $\models \neg B\phi$.

But this principle alone does not yet ground inferences about the beliefs an agent lacks from the beliefs they have. For that, we have to ask when inconsistencies between multiple beliefs count as blatant. Cherniak suggests *contradictories* are examples of this kind (i.e. a proposition and its negation). That is, if a minimally rational person believes ϕ , they do not believe $\neg\phi$ as well (p. 16). As with conjunction elimination, Cherniak's intuition about contradictories is widely shared (in particular, the notions of minimally rational belief of Vardi 1986, Jago 2013 and Solaki et al. 2019 underwrite this principle.) So plausibly, *Minimal Consistency* should also entail the following principle:

Avoidance of Contradictories. Rational agents do not believe contradictories: $B\phi \models \neg B\neg\phi$.

The Puzzle about Consistency arises from the tension between this conception of minimal consistency and another venerable principle of doxastic logic:

Ideal Adjunction. If a rational agent believes some things, they also believe their conjunction. That is: $B\phi, B\psi \models B(\phi \wedge \psi)$.

If we assume *Ideal Adjunction*, then *Ideal Consistency* is equivalent to *Avoidance of Contradictions*. For suppose an agent believes the inconsistent propositions $\phi_1, \phi_2, \dots, \phi_n$. Then given *Ideal Adjunction* they would also believe $(\phi_1 \wedge \phi_2 \wedge \dots \wedge \phi_n)$, which is a contradiction.⁷ Thus accepting *Ideal Adjunction* makes it impossible to formulate a *Minimal Consistency* condition that is intermediate in strength between *Ideal Consistency* and *Avoidance of Contradictions*. In Section I we saw that *Intensionality* collapses the distinction between *Ideal* and *Minimal Closure*, and the distinction between proximate and remote consequences. In a similar way, *Ideal Adjunction* collapses the distinction between *Ideal* and *Minimal Consistency*, and the distinction between blatant and hidden inconsistencies.

Consequently, just as the theoretical demand for *Minimal Closure* casts doubt on *Intensionality*, so the need for *Minimal Consistency* casts doubt on *Ideal Adjunction*.⁸ This doubt is reinforced by the fact that cases of conflicting beliefs typically make for intuitive counterexamples to *Ideal Adjunction*. For instance, suppose I believe that Ann will come to the party even though I know that Tom was also

⁷ For simplicity, I'm ignoring the case of infinitary inconsistencies, as well as infinitary conjunctions.

invited and that Ann avoids Tom like the plague. If someone were to point this out, I would revise my belief that Ann will come. But as things stand, I have the belief that Ann will come, and the belief that Tom will come, but I intuitively lack the belief that Ann and Tom will both come.⁸

A popular response to the Puzzle about Consistency is the fragmentation theory of belief (Lewis 1982, Stalnaker 1984, 1999, §6 of Fagin and Halpern 1988, Egan 2008, Greco 2015, Pérez Carballo 2016, Borgoni, Kindermann, and Onofri 2021, Yalcin 2018, 2021, Elga and Rayo 2021, 2022). According to this view, *Ideal Adjunction* is false because our beliefs are divided into distinct, compartmentalised belief systems. Each fragment is individually consistent, and the conjunctions of beliefs within a single fragment are believed. But if the belief that ϕ is part of one fragment and the belief that ψ is part of another, then the belief $(\phi \wedge \psi)$ need not be part of any fragment.

The fragmentation view accepts *Avoidance of Contradictions* while denying *Ideal Consistency*. So it has a *Minimal Consistency* condition of sorts. But I think this condition is too weak to make for a satisfactory solution to the puzzle. In particular it does not do justice to Cherniak's idea that what we know about the beliefs someone *has* should tell us something about the beliefs they *lack*. By design, the fragmentation view does not sustain any such inferences. The beliefs in different fragments are compartmentalised and do not constrain one another. So switch fragments, and all bets are off. In particular, a fragmentation account of rational belief does not sustain *Avoidance of Contradictories*: a fragmented agent may believe ϕ and simultaneously believe $\neg\phi$ as part of some other fragment. So a truly fragmented agent is not minimally rational in Cherniak's sense.⁹

One could add a stipulation that belief fragments are to be pairwise consistent. That formal fix does give you *Avoidance of Contradictories*, but it seems rather ad hoc. The move also goes against the spirit of the fragmentation view: how are the fragments supposed to stay pairwise consistent if they are compartmentalised? Furthermore, once you allow such add-on stipulations, we get perhaps more freedom than we want. Why stop at pairwise consistency, for instance? What about a constraint that any four fragments need to be mutually consistent? Or any five? Again, a more satisfactory solution to the Puzzle about Consistency would yield insight into minimal rationality, and provide some guidance on how to explicate the intuitive notion of a blatant inconsistency.

⁸ The case against *Ideal Adjunction* is admittedly less clear-cut than the case against *Intensionality*, because the status of contradictions as blatant inconsistencies is intuitively less secure than the status of conjuncts as direct consequences. In particular, you could have a view of blatant inconsistency that rules in some contradictions but not others. Such an account could combine a commitment to *Ideal Adjunction* with a non-trivial *Minimal Consistency* constraint that entails *Avoidance of Contradictories*, say, but not *Avoidance of Contradictions*. But that way out of the puzzle is not available if you endorse *Ideal Closure*, as many of the fragmentation theorists cited below do. For given *Ideal Adjunction* and *Ideal Closure*, any agent with inconsistent beliefs believes every proposition. That makes *Ideal Consistency* equivalent to *Avoidance of Contradictories* and to any other non-trivial consistency constraint you might come up with.

⁹ This relates to a deeper worry about the idea of using fragmentation to capture minimal rationality. The choices of a fragmented agent are guided by different fragments on different occasions (Elga and Rayo 2019). Even if their beliefs are assumed to stay fixed, the beliefs that guide an agent now are not guaranteed to be in effect in five minutes: they might switch fragments in the interval. For this reason, fragmentation threatens to undermine the very coherence that minimal rationality is specifically intended to capture (Norby 2014; Hoek 2019, §2.5).

III. Question-Directed Beliefs

In this section I give an outline of a simple new theory of static minimal rationality that addresses our Puzzles about Closure and Consistency, yielding precise notions of direct consequence and blatant inconsistency. The starting point for my account is a view of cognitive content that has been defended by Seth Yalcin (2011, 2018) and many others on a range of philosophical, linguistic and psychological grounds: namely the view that the objects of belief are answers to specific questions, instead of undirected pieces of information about the world.¹⁰ According to that view, believing that *Paul is going to Paris next week* in answer to the question *When is Paul going to Paris* is not the same thing as believing that *Paul is going to Paris next week* in answer to the question *Where is Paul going next week*. These beliefs have the same truth conditions, but answer different questions. Likewise, the belief that *Either emus can fly or they can't* is distinct from the intensionally equivalent belief that *Either it is snowing or it isn't* because those beliefs answer different questions: *Can emus fly* versus *Is it snowing*. In belief reports, these distinctions are sometimes marked using word choice or focus.

To distinguish those hyperintensional belief contents from other kinds of propositions, I will call them *quizpositions*, short for question-directed propositions. We will model questions as partitions of logical space (in the tradition of Hamblin 1958, Lewis 1988, Groenendijk and Stokhof 1997):

A (*partition*) *question* Q is a partition of logical space Ω . The cells $q \in Q$ of this partition are called *Q-cells*. When two worlds w and v share a Q -cell, we write $w \sim_Q v$. Any set of Q -cells $A \subseteq Q$ is an *answer* to Q . (a)

A *question-directed proposition* or *quizposition* is an ordered pair $\langle Q, A \rangle$, also denoted A^Q , consisting of the partition question Q that A^Q is said to be *about*, and some answer $A \subseteq Q$. The quizposition A^Q is *true* at a Q -cell q if and only if $q \in A$, and it is *true* at a world w if and only if $w \in \bigcup A$. (b)

The singleton sets $\{q\} \subseteq Q$ are *complete* answers to the question Q ; all other non-empty subsets of Q represent *partial* answers. For instance, *There are fewer than ten people in the room* is a partial answer to the question *How many people are there in the room*. Any question Q has a tautologous answer Q and an absurd answer \emptyset ; the corresponding quizpositions are written Q^Q and \perp^Q .

How does the move from intensional propositions to quizpositions help us formulate a principled notion of minimal rationality? Let me briefly sketch an answer. The intensional, sets-of-worlds account of belief explicates belief in terms of our ability to rule out ways the world might be. This picture ignores the role of another, prior cognitive ability: namely the ability to distinguish between various possibilities in the first place. In taking beliefs to be sets of possible *worlds*, the traditional account in some sense presupposes that believers individuate possibilities maximally finely, and have already made every distinction there is to make.

¹⁰ Defences include Dretske 1970, Schaffer 2007, Egré and Bonnay 2012, Blaauw 2013, Koralus and Mascarenhas 2013, 2018, Yablo 2014 (Ch. 7), Ciardelli and Roelofsen 2015, Fritz and Lederman 2015, Pérez Carballo 2016, Friedman 2017, Bledin and Rawlins 2018, Drucker 2020 and Holguín 2022. The main development from Yalcin's question-sensitive theory of belief to the account below is the addition of coherence constraints between agents' answers to different questions. Being a fragmentation theorist, Yalcin posits no such constraints (see also Yalcin 2021). Another difference is that for Yalcin, the contents of beliefs are still sets of worlds, and not question-directed propositions. In Yalcin 2018, he draws explicit attention to the fact that his account still validates *Intensionality*, which he labels "*Closure under Necessarily Equivalence*". Yalcin does in effect address the Puzzle about Closure. But his solution is to accept *Intensionality* and to reject *Closure under Conjunction Elimination*: on Yalcin's account, one can believe $A \wedge B$ in answer to $Q \wedge R$ without believing A in answer to Q .

The question-directed view of belief eliminates that idealisation. It holds that, prior to forming a contingent belief, a subject must conceptualise the relevant possibilities. If a question Q has four cells, say, one can come to believe the quizposition A^Q only after first distinguishing those four possibilities. Forming a view on a given question is itself a substantive cognitive achievement, which takes mental resources and requires developing a connection or sensitivity to the corresponding aspect of what the world is like. This makes it natural to assume that doxastic rationality does not require us to ask particular questions, any more than it requires us to give particular answers: it can only constrain how our views on questions we do grasp should cohere with one another. By following that line of thought, a picture will emerge of a minimally rational believer whose deductive achievements and limitations can be systematically understood in terms of the possibilities they have and have not distinguished, and the questions they have and have not asked.

Question Mereology

In natural language, interrogatives can be conjoined in the same way as declarative sentences. This naturally gives rise to a notion of question *conjunction* and the related notion of question *parthood*, both of which will be important for my account. Consider for example the conjunctive question *How many stars are there and how many planets are there*. One complete answer to that question is *exactly twenty-five stars and three planets*, and in general any complete answer to the conjunctive question is a conjunction of a complete answer to *How many stars are there* and a complete answer to *How many planets are there*. Generalising that pattern, we get the following definition:

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

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

QR is the partition such that $w \sim_{QR} v$ if and only if $w \sim_Q v$ and $w \sim_R v$. (c)

Not every partial answer to QR is a conjunction of an answer to Q and an answer to R : for instance, one partial answer to *How many stars and planets there are* is *There are more planets than stars*.

Note that Q and R are both *coarser* partitions than QR , in the sense that each Q -cell and each R -cell is a union of smaller QR -cells (in fact, QR is just the coarsest common refinement of Q and R). The notion of a question part is a generalisation of the relation question conjuncts bear to their conjunction:

One question Q *contains* (or is *at least as big as*) another question R if and only if every R -cell is a union of Q -cells. We say 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 $w \sim_Q v$. (d)

Note that Q contains R if and only if $QR = Q$. Big questions draw more distinctions between possibilities than the questions they contain. Less abstractly, one question is part of another if it has to be resolved to get a complete answer to the bigger question. For instance, *What month is it* is part of *What date is it* and *What street does Jess live on* is part of *What is Jess's address*. The trivial question $\{\Omega\}$, drawing no distinctions at all, is part of every question.

Figure 1 below provides a visual illustration. Each square represents a question: a partition of logical space. The question at the top makes more distinctions than those displayed below it, and so it *contains* those questions as parts. In fact, since it is the smallest (most coarse-grained) question to contain both of them as parts, it is the conjunction of the smaller questions.

A conjunction of question parts is itself always a part. Hence the common parts of any two given questions are closed under conjunction, so that there is always a greatest common part:

The *overlap* of two questions Q and R is the biggest question that is part of Q and also part of R. Two questions *overlap* if and only if their overlap is not $\{\Omega\}$. Otherwise they are *disjoint*. (e)

For instance, the question *What are the capitals of Europe* overlaps *What are the capitals of Asia*, and their overlap is the question *What are the capitals of Turkey and Russia* (assuming for the sake of the example that it is not contingent what the countries of Asia and Europe are). Any answer to the latter question is a partial answer to both of the bigger questions. The more (partial) answers two questions have in common, the more they overlap, and two questions do not overlap at all if they do not have contingent partial answers in common.




FIGURE 1: QUESTION PARTHOOD AND QUESTION CONJUNCTION

Minimally Rational Answers

An answer to a big question also answers all of its parts. A view on *What date it is* says something about *What month it is* and something about *What day of the month it is*. So given the view that beliefs *are* answers to questions, it is natural to expect that an agents' beliefs about a question should harmonise with their beliefs about the parts:

Harmonic Parts. If a rational agent has beliefs about Q, then they have *matching* beliefs about every part R of Q: that is to say, they believe all and only those quizpositions about R that are entailed by their beliefs about Q.

This is a pretty intuitive constraint. An agent's view about the whole includes and reflects their views about the parts. So if you believe *It is the 13th of March* (in answer to *What date is it*), plausibly you also believe that *it is March* (in answer to *What month is it*). And if I am unsure whether *It is the 13th of March or April*, it intuitively follows that I must also be unsure whether *It is March or April*.

Recall that the trivial question $\{\Omega\}$ is part of every other question. So it follows from *Harmonic Parts* that if an agent believed the trivial absurdity $\perp^{\{\Omega\}}$, they would have inconsistent beliefs about every question to which they had an answer. To exclude that possibility, we will assume that minimally rational agents cannot be in such a state:

Non-Absurdity. Rational agents do not believe $\perp^{\{\Omega\}}$.

That completes the account of static minimal rationality I want to propose: at a given time, an agent

has minimally rational beliefs if they satisfy *Harmonic Parts* and *Non-Absurdity*. In Section IV below, I show how this account yields attractive *Minimal Closure* and *Minimal Consistency* conditions.

Harmonic Parts and *Non-Absurdity* are limited constraints that are in principle attainable by a finite reasoner whose beliefs concern questions with finitely many cells. They only require the agent to integrate a given belief with beliefs related to it, rather than with all other beliefs. I will seek to make it plausible that, besides being attainable in principle, ordinary agents for the most part also meet these constraints in practice (by “ordinary agents” I have in mind normal adult human beings, say).

At the same time, *Harmonic Parts* and *Non-Absurdity* impose a greater amount of coherence on an agent’s beliefs than meets the eye. Let me explain why. An ordinary agent presumably has views on a wide variety of questions. Wherever those questions overlap, *Harmonic Parts* directly constrains the relationship between an agent’s views on those questions: if a minimally rational agent has beliefs on questions Q and R which share a part S, then their beliefs on Q and R must entail all and only the same answers to S. For instance, minimally rational views about *How old Alice and Bob are* and about *How old Bob and Carmen are* must always coincide on the issue of Bob’s age.

The theory also forges indirect links between the agent’s views on disjoint questions. The reason is that those views may be connected through one or more daisy-chains of background beliefs about overlapping questions. Figure 2 illustrates this. The questions on the top left and top right are disjoint: they do not make any of the same distinctions. Nonetheless, they each overlap with the question in the middle: the shared parts are shown underneath. The questions addressed by an agent’s beliefs will usually form a complex mereological structure, naturally described as a web: hence the *web of questions*. Broadly speaking, an agent’s views on some questions become better integrated and more coherent the more connected those questions are within the agent’s web. And while some questions may be more interconnected than others, few if any will be entirely isolated from the overall web.




FIGURE 2: A DAISY-CHAIN OF OVERLAPPING QUESTIONS

In this way, *Harmonic Parts* and *Non-Absurdity* impose substantial yet attainable constraints on minimally rational beliefs. Thus the present account aims to do justice to Cherniak’s observation that “the assumption that the agent can make quite complex inferences from his beliefs is crucial to our pretheoretical attributions of psychological states in everyday situations” (p. 28).¹¹

IV. Direct Consequences and Blatant Contradictions

We now have a simple theory of static minimal rationality. This section explains how this theory addresses the Puzzles about Closure and Consistency. In particular, I isolate the *Minimal Closure*, *Minimal Adjunction* and *Minimal Consistency* conditions entailed by the theory. This yields sharply defined notions of a *direct* consequence (the sort of consequence any minimally rational agent can see), and of a *blatant* inconsistency (the sort of inconsistency any minimally rational agent will avoid).

Minimal Closure

Section I hypothesised that minimally rational agents believe the *direct* consequences of their beliefs, including the conjuncts of every conjunction they believe. To appreciate what that implies in the present context, we first need to define quizposition conjunction. Given our definition of question conjunction, there is only one sensible way to do this:

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 $\langle QR, AB \rangle$. (f)

A quizposition conjunction makes just enough distinctions between possible worlds to make every distinctions that its conjuncts make, and rules out just enough possibilities to rule out every possibility that its conjuncts rule out.

This yields a notion of propositional (quizpositional) *parthood* comparable to that of Yablo 2014 and Fine 2017. Just as question parthood is the relation that question conjuncts bear to a conjunctive question, quizposition parthood is the relation quizposition conjuncts bear to their conjunction. That is, one quizposition is part of another if it makes fewer distinctions and rules out fewer possibilities:

A quizposition A^Q *contains* a quizposition B^R , or B^R is *part of* A^Q , if and only if Q contains R and A entails B (that is, $\bigcup A \subseteq \bigcup B$). (g)

As in the case of questions, one quizposition contains another just in case the conjunction is equal to the whole. That is to say, A^Q contains B^R if and only if $AB^{QR} = A^Q$. Not every part is an explicit conjunct. For instance, the quizposition *Gold is a soft yellow metal* (in answer to *What are the properties of gold*) contains the quizposition *Gold is yellow* (in answer to *What is the colour gold*). And *Fred's phone number starts with a four*, in answer to *What is the first digit of Fred's phone number*, is part of *Fred's phone number is 49753*, in answer to *What is Fred's phone number*.

If the objects of belief are intensional, $B\phi \models B(\phi \wedge \psi)$ just in case ϕ entails ψ : that is why *Closure under Conjunction Elimination* collapses into *Ideal Closure* in an intensional context. But on the quizpositional account of belief, $B\phi \models B(\phi \wedge \psi)$ just in case ϕ and $(\phi \wedge \psi)$ express the same quizposition, which happens just in case ϕ contains ψ as a part. So in the present setting, *Closure under Conjunction Elimination* yields the following principle:

Closure under Parthood. Whenever a rational agent believes something, they believe all of its parts. So if $\phi \leq \psi$, then $B\phi \models B\psi$.

¹¹ Still, are these constraints strong enough to sustain our ordinary belief-based explanations and predictions of behaviour and decision-making, the way Cherniak envisioned? I think so, but admittedly the matter requires more attention than I can give it here. I say a good deal more in Hoek 2022, which approaches the problem of logical omniscience from a decision-theoretic angle.

Here the notation “ $\phi \leq \psi$ ” abbreviates “the quizposition expressed by ϕ contains the quizposition expressed by ψ ”.¹² *Closure under Parthood* is precisely the *Minimal Closure* principle we get from the present theory: it is a straightforward consequence of *Harmonic Parts*. Yablo 2014 (Ch. 7) and Hawke 2016 have defended an analogous closure principle for knowledge.

Intuitively, *Closure under Parthood* covers precisely the sort of automatic inferences we want a *Minimal Closure* condition to capture. For instance, it is immediately plausible that believing that *gold is a soft yellow metal* implies believing *gold is yellow*. Likewise, believing that *Mary lives on 15 Baker Street* seems to involve believing that *Mary lives on Baker Street*. Or suppose I ask myself *What time is it*, and take a look at my watch: as I find out *It's three forty-five*, I also acquire the belief that it is *not yet four o'clock*.

More tellingly still, as Yablo points out, *Closure under Parthood* yields intuitively compelling explanations for the failures of *Ideal Closure*. For instance, believing that *The wall is blue* does not entail believing *Either the wall is blue or there is something wrong with the lights*, because the latter belief involves consideration of a bigger question, *What colour is the wall and what is the condition of the lights*. Believing *Zed is a zebra* does not entail believing *Zed is not a cleverly disguised mule*, because only the latter belief is about disguise. King William believed that *England could avoid war with France* but not that *England could avoid nuclear war with France*, because the latter belief answers a question no one in the eighteenth century was even able to pose (cf. Yalcin 2011, §8). And we can begin to see, faintly, why believing the second-order Peano axioms does not entail believing that there are infinitely many primes. Issues about mathematical necessity aside, none of those axioms say anything about primes, or about how many there are (for more on this see Hoek 2019, Ch. 5).

Minimal Adjunction

Before moving on to the Puzzle about Consistency, it will be helpful to discuss where the present theory leaves closure under adjunction. Since any question Q is part of itself, it follows from *Harmonic Parts* that an agent believes all and only those quizpositions about Q that are entailed by their beliefs about Q . In particular, that implies that minimally rational agents believe the conjunction of all their beliefs about a given question. So:

Internal Adjunction. A rational agent believes the conjunction of all their beliefs about any particular question. That is, if $\phi \approx \psi$, then $B\phi, B\psi \models B(\phi \wedge \psi)$.

Here the notation “ $\phi \approx \psi$ ” abbreviates “the quizpositions expressed by ϕ and ψ are about the same question”. Let us say an agent’s *view* about a given question Q is the conjunction of all their beliefs about Q . Together, *Internal Adjunction* and *Closure under Parthood* tell us that a rational agent with beliefs about Q believes all and only those quizpositions about Q that their view on Q entails.

Suppose an agent has a view V^Q on Q . Now consider any part R of Q . *Closure under Parthood* says that the agent believes every quizposition about R that V^Q entails. The other half of *Harmonic Parts* is the

¹² This gloss of “ \leq ” corresponds to an understanding of the statement “if $\phi \leq \psi$, then $B\phi \models B\psi$ ” as a schema in which ϕ and ψ range over English declarative sentences, and “ B ” abbreviates “ α believes that” (for a semantics of “believe” that validates this principle and the other principles of minimal rationality advocated here, see appendix). Like *Ideal Closure*, *Closure under Parthood* can also be understood as a property of a logic of belief. In that context, “ \leq ” expresses the relation of *analytic entailment*, or parthood under every interpretation (cf. Fine 2016). That relation can be characterised axiomatically, as in Goodman 2019.

converse: the minimally rational agent believes *only* those quizpositions about R that V^Q entails. This is encoded in the following *Minimal Adjunction* principle:

Partial Adjunction. Rational agents' beliefs on any part of a question are incorporated into their view about the whole question. That is, if $\phi \approx (\phi \wedge \psi)$, then $B\phi, B\psi \models B(\phi \wedge \psi)$.

Internal Adjunction is a consequence of *Partial Adjunction*.

For an intuitive motivation of *Partial Adjunction*, consider situations where a rational agent is unsure about the conjunction in question. In such cases, it intuitively follows that they must be unsure about the relevant conjunct too. For example, suppose you believe Beth's house number is 22, but are unsure whether she lives on 22 *Broad Street* or 22 *High Street*. It seems to follow that you must be unsure about *What street Beth lives on*. Contrapositively, if you firmly believed *Beth lives on Broad Street*, that belief would settle your view on *What Beth's address is*. Likewise, if you are unsure whether John's phone number is 76453 or 86453, you are unsure about the first digit. If you believe *Bismuth is either a hard, reddish metal or a soft, blueish metal* but are unsure which it is, then you are uncertain about both the hardness and the colour of bismuth. And so on. *Partial Adjunction* makes systematic sense of these intuitions. As with *Closure under Parthood*, this constraint is not just attainable in principle. It is at least *prima facie* plausible that ordinary, finite agents for the most part attain it.




FIGURE 3: THE VIEW ABOUT THE WHOLE DETERMINES THE VIEWS ABOUT THE PARTS

Taken together, *Closure under Parthood* and *Partial Adjunction* are equivalent to *Harmonic Parts*: when R is part of Q, a rational agent with a view on Q believes all and only those quizpositions about R that their view on Q entails. So on the present theory, a minimally rational agent's view on Q completely determines their view on every Q-part R. More specifically, if V^Q is their view on Q, the agent's view on R is the strongest quizposition about R entailed by V^Q .

Figure 3 above illustrates the principle. Each square represents a quizposition. The black lines make a partition representing a question (as before): this is the inquisitive component of the quizposition. The colouring represents its informational component: light grey marks the cells where the quizposition is true and a darker grey the cells where it false. The three quizpositions displayed at the bottom are part of the quizposition at the top: they are weaker, and make fewer distinctions. More specifically, they each rule out all and only those answers ruled out by the view at the top. So if the top quizposition represents a minimally rational agent's view on some big question, then the quizpositions displayed underneath are the views this agent must hold about its component questions.

Minimal Consistency

Now we are ready to address the Puzzle about Consistency. Let me start by checking that from *Harmonic Parts* and *Non-Absurdity*, we can recover *Avoidance of Contradictions* and *Avoidance of Contradictories*. To evaluate the latter principle, we need to define quizposition negation:

The *negation* of a quizposition A^Q , written $\neg A^Q$, is the quizposition $\langle Q, Q \setminus A \rangle$. (h)

From *Harmonic Parts*, we get that agents believe both A^Q and $\neg A^Q$ just in case they believe \perp^Q . Now since $\{\Omega\}$ is part of any question Q , $\perp^{\{\Omega\}}$ is always part of \perp^Q , and thus *Non-Absurdity* rules out belief in contradictions and also belief in contradictories.

More generally, the theory has it that an agent cannot have inconsistent beliefs about a particular question Q , since those would always adjoin to \perp^Q . This accords pretty well with our pre-theoretical intuitions. If I am at all rational, I cannot both be convinced that *my coat is red all over* and also that *it is blue all over*. If I believe *I left the keys in the car*, I do not think *the keys are in my pocket*. If I look at the clock on the wall to discover *it is 3 o'clock*, I am forced to discard the belief that *it is 2:30*. And so on. The pattern extends beyond pairwise consistency. It would be odd if you were certain that *Montesquieu was either a novelist, an architect or a banjo player*, while also being convinced that *he was definitely not a novelist*, that *he was definitely not an architect* and that *he was definitely not a banjo player*.

So question-internal inconsistencies are blatant inconsistencies. Are all blatant inconsistencies question-internal? Given *Closure under Parthood*, the answer must be “no”. For instance, consider the belief *Jill and Jack are over twenty-one* and the belief *Jane and Jill are under eighteen*. These quizpositions concern different questions: *How old are Jill and Jack* and *How old are Jane and Jill* respectively. And yet these two beliefs are blatantly inconsistent, both intuitively speaking and according to the theory. For the two questions overlap on the question *How old is Jill*; and it is part of the first quizposition that *Jill is over twenty-one* and part of the second quizposition that *Jill is under eighteen*; and those parts make a question-internal inconsistency. So you cannot believe both of the wholes either.

Generalising from this, minimally rational beliefs about overlapping questions cannot contradict each other on the overlap. That gives us the following *Minimal Consistency* constraint:

Partial Consistency. A rational agent’s beliefs agree on any question. If $\phi_1 \approx \phi_2 \approx \dots \approx \phi_n$, $\phi_1, \phi_2, \dots, \phi_n \models \perp$, and for any i , $\psi_i \leq \phi_i$, then $B\psi_1, B\psi_2, \dots, B\psi_{n-1} \models \neg B\psi_n$.

So a *blatant* inconsistency arises whenever overlapping quizpositions contradict one another on a question in the overlap. When some quizpositions are inconsistent but not blatantly so, we can speak of an *opaque* inconsistency. The two quizpositions at the top of Figure 4 below are opaquely inconsistent. They are inconsistent because there is no point in logical space at which both are true. And they overlap, because they make some of the same distinctions. But nevertheless they agree about the answer to the overlapping part, displayed underneath. So these beliefs are not blatantly inconsistent, and the present theory admits such combinations of views.

Let’s look at a more concrete example. Here are three opaquely inconsistent quizpositions:

- 1) *I have Wednesday free* in answer to *What do I have to do on Wednesday?*
- 2) *I have to go to Sarah’s birthday party* in answer to *Do I have to go to Sarah’s birthday?*
- 3) *Sarah’s birthday party is on Wednesday* in answer to *When is Sarah’s birthday party?*

Quizpositions (1-3) are about non-overlapping questions, so they are opaquely inconsistent. Now if you conjoin all three with *Today is Monday*, the three resulting conjunctions (1*-3*) overlap, but they are still opaquely inconsistent. That's because drawing attention to what day it is does nothing to reveal the inconsistency between (1-3). While (1*-3*) overlap on the question *What day is it today*, they all agree on the overlapping part: the inconsistency lies elsewhere.




FIGURE 4: OPAQUE INCONSISTENCY BETWEEN VIEWS ON OVERLAPPING QUESTIONS

By contrast if we conjoin (1) and (3), and also conjoin (2) and (3), then we do get a blatant inconsistency, both intuitively speaking and according to the theory:

- 4) *I have Wednesday free and Sarah's birthday party is on Wednesday*
- 5) *Sarah's birthday party is on Wednesday and I have to go to Sarah's birthday party.*

Quizpositions (4) and (5) overlap on the polar question *Do I have to go to Sarah's birthday party on Wednesday or not?*, and directly contradict each other there, answering *No* and *Yes* respectively. So *Partial Consistency* rules out this pair of views. In this way, deductive inferences may reveal inconsistencies: with two adjunctions, we went from an opaque to a blatant contradiction.

Here is a more interesting example of an opaque inconsistency between overlapping questions. Let H, N and C be the questions *How tall is Hob*, *How tall is Nob* and *How tall is Cob* respectively. Then these three quizpositions are opaquely inconsistent, even though they all overlap one another:

- 6) *Hob is taller than Nob* in answer to HN
- 7) *Nob is taller than Cob* in answer to NC
- 8) *Cob is taller than Hob* in answer to CH

This example is reminiscent of the opaque inconsistency observed in a Penrose triangle or an Escher print (Figure 5). Such images initially appear consistent, since any sufficiently small region provides a consistent representation of part of the world, just as (6), (7) and (8) are each consistent answers to their respective questions. In both cases, these consistent representations are not compartmentalised. On the contrary, they all overlap with one another, and you cannot draw sharp boundaries between them. Yet they do not add up to a consistent representation of how things are.

This is clearly very different from the account of inconsistent belief given by fragmentation theorists. On that account, inconsistency arises when discrete, independent doxastic representations offer directly contradictory views of the world. It may be that this kind of belief fragmentation occurs in people who have split brains or dissociative personalities. But those cases are pathological, and




FIGURE 5: A PENROSE TRIANGLE AND DETAIL OF M.C. ESCHER'S "WATERFALL"

plausibly fall short even of minimal rationality. I contend that the present account paints a more realistic picture of inconsistent belief as it arises in non-pathological subjects, as a consequence of ordinary failures of logical omniscience.

V. A Puzzle About Deductive Inference

Taken together, the principles of minimal rationality defended above yield a new definition of rational belief states. The standard definition of a rational belief state, notably exemplified by Jaakko Hintikka's (1992) model of rational belief, runs like this:

An *ideal belief state* is a non-empty set \mathbf{I} of intensional propositions (that is, sets of worlds) subject to these three conditions:

- ▶ *Ideal Closure*: If p entails q and $p \in \mathbf{I}$, then $q \in \mathbf{I}$.
- ▶ *Ideal Adjunction*: If $p, q \in \mathbf{I}$, then $(p \wedge q) \in \mathbf{I}$.
- ▶ *Ideal Consistency*: $\perp \notin \mathbf{I}$. (i)

This definition uses a simplified statement of *Ideal Consistency*: conditional on *Ideal Adjunction*, the consistency of \mathbf{I} is equivalent to $\perp \notin \mathbf{I}$.

The analogous definition for a minimally rational belief state replaces each clause in (i) with its minimal analogue:

A *minimally rational belief state* is a non-empty set \mathbf{B} of quizpositions subject to the following three conditions:

- ▶ *Closure under Parthood*: If A^Q contains B^R and $A^Q \in \mathbf{B}$, then $B^R \in \mathbf{B}$
- ▶ *Partial Adjunction*: If Q contains R , and $A^Q, B^R \in \mathbf{B}$, then $AB^Q \in \mathbf{B}$
- ▶ *Partial Consistency*: $\perp^{(\Omega)} \notin \mathbf{B}$. (j)

Again, this definition uses a simplified consistency clause. Conditional on the first two clauses, *Partial Consistency* is equivalent to $\perp^{(\Omega)} \notin \mathbf{B}$ (*Non-Absurdity*).

Perhaps the most exciting aspect of this new account of belief states is the way it allows us to theorise systematically about deductive inferences. A deductive inference is the formation of new a belief on the basis of extant beliefs that entail it. Ideal belief states preclude the very possibility of such transitions, because they are already deductively closed. *Intensionality* is by itself compatible with

failures of deductive closure, as witnessed by the fragmentation view discussed above, as well as the neighbourhood models of Montague 1970 and Scott 1970. But any intensional view individuates inferences very coarsely, and that makes it difficult to form a realistic picture of deductive reasoning on an intensional basis. To bring that out, this section turns to another question raised by Cherniak's account: what makes some deductive inferences more difficult than others?

As I mentioned at the beginning of the chapter, Cherniak holds that minimal rationality has a dynamic aspect. For Cherniak, minimal rationality not only requires that you have fairly cogent beliefs, but also that you make sensible inferences from those beliefs when the need arises. Of course there is a limit on what can reasonably be demanded. Cherniak holds that the complexity of the deductive inferences that a subject can reliably be expected to make depends on the time and cognitive resources that are available for thinking about the issue at hand.

Common sense offers some guidance about which deductions require greater cognitive resources: the inference from the truth of the Peano axioms to the truth of Fermat's last theorem, for instance, was extremely difficult. It was much harder than the inference from the clue entries of a simple Sudoku to its solution. Solving a Sudoku, in turn, is harder than performing an instance of universal instantiation. However, to render this aspect of Cherniak's view predictive, we need something more systematic than those case-by-case judgments. We need some independent handle on what makes particular inferences more or less difficult. Ideally, this would allow us to order deductive inferences by difficulty in a hierarchy, and this is what Cherniak envisaged.

However, *Intensionality* makes it very difficult to make sense of Cherniak's vision. In the case of single-premise inferences, we have already seen why: modulo *Intensionality*, every valid single-premise inference is an instance of Conjunction Elimination. That would make Conjunction Elimination the hardest single-premise inference rule there is, since it subsumes all the others. Putting it differently, given *Intensionality*, every single-premise inference is as easy as a Conjunction Elimination. That leaves no space for any interesting hierarchy in the domain of single-premise inferences.

An analogous issue arises for multi-premise inferences. Assuming *Intensionality*, every deductively valid inference is an instance of the following inference rule:¹³

Recombination. $(\phi_1 \wedge \psi_1), (\phi_2 \wedge \psi_2), \dots, (\phi_n \wedge \psi_n) \quad \therefore \quad (\phi_1 \wedge \phi_2 \wedge \dots \wedge \phi_n)$

Intuitively, the Recombination rule looks like a trivial inference: it simply re-asserts some conjuncts in the premises. Nonetheless, given *Intensionality*, Recombination encompasses everything from the humblest modus ponens to the highest flights of human reason. You could say *Intensionality* puts a ceiling on how difficult a deductive inference can be, and that ceiling looks uncomfortably low. On the face of it, *Intensionality* forces the bizarre conclusion that every deductive argument ever made is really only a repetition of some judiciously selected premises. If one accepts this, it is hard to make sense of Cherniak's hierarchy in the case of multi-premise inferences, too. The Puzzle about Deductive

¹³ *Proof.* To show: if $\alpha_1, \alpha_2, \dots, \alpha_n \models \beta$, then the inference from $\alpha_1, \alpha_2, \dots, \alpha_n$ to β is an instance of *Recombination*, modulo *Intensionality*. To see this, simply substitute $\phi_i = (\beta \vee \alpha_i)$ and $\psi_i = \alpha_i$. For note $\alpha_i \models (\phi_i \wedge \psi_i)$, and because of distributivity, we have $\beta \models \beta \vee (\alpha_1 \wedge \alpha_2 \wedge \dots \wedge \alpha_n) \models (\beta \vee \alpha_1) \wedge (\beta \vee \alpha_2) \wedge \dots \wedge (\beta \vee \alpha_n) = (\phi_1 \wedge \phi_2 \wedge \dots \wedge \phi_n)$.

Inference is the question of how we are to resolve the resulting tension between *Intensionality* and the undeniable fact that deductive accomplishments often require considerable effort.

Arguably, the intensionalist's prospects for explaining the effort involved in multi-premise inferences are a little better than for single-premise inferences. That is because the Recombination rule involves *adjunction* (conjunction introduction) as well as conjunction elimination. And amongst advocates of fragmentation theories, there is a tradition holding that conjoining separated beliefs into a single belief is precisely what deductive inference amounts to: "There may be propositions which I would believe if I put together my separate [fragments] of belief, but which, as things stand, hold in none of them. These are the propositions that may be discovered by a purely deductive inquiry." (Stalnaker 1974, p. 85) If that is right, and these adjunctions are what takes cognitive effort, then Recombinations presumably get more difficult as they involve more adjunctions. Could this be the hierarchy Cherniak envisioned? Could adjunction be the intensionalist's Archimedean point, the one source of friction that will put a distance between obvious and remote consequences?

Probably not. Deriving the commutativity of multiplication from the second-order Peano axioms requires eight adjunctions, since there are nine axioms. Deriving Fermat's last theorem? Again, eight. The Goldbach Conjecture? Eight again, assuming it is true. Simply counting the number of adjunctions made is clearly no guide to the impressiveness of a deductive accomplishment. Perhaps we can instead distinguish easy and difficult adjunctions. According to Stalnaker, adjunction "may require only a routine calculation, or it may be a challenging and creative intellectual task." (1984, p. 84) Still, there is apparently no way of anticipating, in any given case, which it is. To get at the Goldbach Conjecture, nine beliefs must be conjoined: (*Peano Axiom 1* \vee *Goldbach Conjecture*), ..., (*Peano Axiom 9* \vee *Goldbach Conjecture*). Which of these adjunctions is it that has baffled the world's greatest mathematicians for three centuries?

The fragmentation theorist's project of reducing all deductive reasoning and information processing to adjunction undeniably has a certain heroic charm. But if we take it seriously, this vision of inference quickly begins to look very implausible (see also Jago 2014, §2.5). Even if the reader does not agree that the intensional view renders the cognitive distinctions we are after more puzzling, I hope I have said enough to persuade them that the intensional view does not *help* dissolve the mystery either. As with the two puzzles about static rationality, what we are really looking for in a solution to this Puzzle about Deduction is a view of belief that makes some progress on the question we started off with. That is, we want a view that yields clarification on what it is that makes deductive inferences challenging. In the remaining two sections, I will try to show how the question-centric view of minimal rationality meets that demand.

VI. Tautological Belief Updates

In this section, we will make a foray into the dynamics of minimally rational belief, by defining belief updates for minimally rational belief states. This yields the notion of a *tautological* belief update, which gives us a natural way to model some deductive inferences. In Section VII, I will then discuss three sources of cognitive difficulty that arise in performing these updates, and relate those observations to experimental findings from psychology.

Belief Updates

The basic dynamic notion in regular doxastic logic is that of an informational update, representing the way an ideally rational subject acquires new beliefs. Given a belief state \mathbf{I} and a new proposition p , the updated state $\mathbf{I} + p$ is the smallest set of propositions closed under entailment and adjunction that has $\mathbf{I} \cup \{ p \}$ as a subset. $\mathbf{I} + p$ is only a belief state when p is consistent with \mathbf{I} . To adopt a belief that is inconsistent with their prior beliefs, an ideal agent would first have to revise their beliefs.

Quizpositional updates can be defined in an exactly analogous way. Call a set of quizpositions *harmonic* just in case it is closed under parthood and partial adjunction (that is, just in case anyone who believed just those quizpositions would satisfy *Harmonic Parts*).

The *update* of a harmonic set of quizpositions \mathbf{B} by a quizposition A^Q , written $\mathbf{B} + A^Q$, is the smallest harmonic set containing $\mathbf{B} \cup \{ A^Q \}$. (k)

Since the set of all quizpositions is harmonic, $\mathbf{B} \cup \{ A^Q \}$ always has a harmonic superset. Since any intersection of harmonic sets is itself harmonic, $\mathbf{B} \cup \{ A^Q \}$ always has a minimal harmonic superset. So $\mathbf{B} + A^Q$ is always well-defined.

As in the ideal case, the result of a quizpositional update is not always a belief state. Since minimal rationality allows for inconsistency, updating by a quizposition that is inconsistent with the anterior state need not be a problem. But if the update fails to preserve *Partial Consistency*, some sort of belief revision will be needed before the update can be performed (cf. Berto 2019). For present purposes, we will set such difficult cases aside, focussing on updates that do preserve *Partial Consistency*.

Tautological Updates

Updating an ideal belief state with a necessary truth leaves the state unaffected. But updating a minimally rational belief state with a necessarily true quizposition Q^Q can yield new beliefs, including new contingent ones. Drawing new distinctions enriches your prior views, bringing them to bear on larger questions with more parts. For instance, you can get from *England will avoid war with France* to the conclusion that *England will avoid nuclear war with France* by drawing a distinction between nuclear and other kinds of war. New questions can also link previously separated views. For instance, to get from *Hob is five foot five* and *Nob is five foot six* to the conclusion *Nob is taller than Hob*, you need to ask *How tall Hob and Nob are*: that question brings both pieces of information together into a single view. Since tautological updates only yield new beliefs that are entailed by the subject's prior beliefs, they make a natural model of deductive inference.

To see how this works, let me go through the two examples just mentioned in more detail. Let E be the polar question *Will England have a war with France or not*. Now suppose Mary's belief state \mathbf{B}_M contains the quizposition A^E , that *England will avoid war with France*. Let F be the tripartite question *Will England have a nuclear war with France or some other kind of war or no war at all*: this question contains E . The tautologous answer to F is F^F , that *Either England will have a nuclear war with France or some other kind of war or no war at all*. If Mary has the ability to distinguish the possibility of nuclear war with France from other kinds of war, she can reason her way from the prior belief state \mathbf{B}_M to the state $\mathbf{B}_M + F^F$. Besides the quizposition A^E , the posterior state $\mathbf{B}_M + F^F$ includes a view on the new question F . Because $\mathbf{B}_M + F^F$ is harmonic, and the question F contains E , the view on F in $\mathbf{B}_M + F^F$ must rule out every

F-possibility incompatible with A^E . In particular, the view rules out the possibility of nuclear war, and so $\mathbf{B}_M + F^F$ is bound to include the quizposition B^F , *England will avoid nuclear war with France*.

For the second example, let H and N stand for the questions *How tall is Hob* and *How tall is Nob*. Cob's belief state \mathbf{B}_C includes the following two quizpositions:

V^H : *Hob is five foot five*

S^N : *Nob is five foot six*

Suppose Cob now forms a view on the conjunctive question HN for the first time: *How tall are Hob and Nob*. Then he transitions from his anterior state \mathbf{B}_C to the belief state $\mathbf{B}_C + \text{HN}^{\text{HN}}$. To preserve harmony, Cob's newly acquired view on HN must exclude the possibilities excluded by V^H , as H is part of HN. By the same token, the view will exclude every possibility excluded by S^N , since N is also part of HN. Thus $\mathbf{B}_C + \text{HN}^{\text{HN}}$ contains VS^{HN} , the conjunction of V^H and S^N . So this tautological update effectively amounts to adjoining of Cob's views on H and N. Now the quizposition

T^{HN} : *Nob is taller than Hob*

is part of VS^{HN} . So because of *Closure under Parthood*, $T^{\text{HN}} \in \mathbf{B}_C + \text{HN}^{\text{HN}}$.

As described at the end of Section III above, the beliefs of a minimally rational agent form a web of interconnected views. For this reason, the effects of a belief update need not be restricted to views on questions that are directly related to the new quizposition. For instance, given suitable background beliefs in \mathbf{B}_C , the realisation that *Nob is taller than Hob* might be accompanied by the realisation that *Gob is taller than Hob* too, which could in turn affect Cob's opinion about *How tall Gob is*, even though the latter question is disjoint from the question HN we updated with.

Figure 6 below illustrates the abstract situation. Here we start out with a belief state containing views on Q_0 , R and S, and then perform a tautological update to refine first of these questions Q_0 to Q_1 . The resulting view $A_1^{Q_1}$ has the same truth conditions as $A_0^{Q_0}$. However, the new question Q_1 that this view addresses overlaps with R, and $A_1^{Q_1}$ rules out some cells in the overlapping part. Consequently, the update strengthens the agent's view on R from B_0^R to B_1^R to preserve harmony. This change in view about R in turn strengthens the agent's view on the question S, which also overlaps with R. Thus the update by $Q_1^{Q_1}$ causes a change in view about S, even though the new question Q_1 does not overlap with S. In the same way, the update could percolate further down the daisy chain, spreading through the agent's web of questions. In this way, a tautological update with $Q_1^{Q_1}$ can in principle affect the agent's view on all kinds of questions that are linked only indirectly to Q_1 .

So on the present model, the acquisition of new tautologous beliefs can lead to all sorts of new contingent beliefs, including beliefs about questions that are pretty remote from the question the new tautology directly addresses. We can think of tautological updates as modelling what happens when an agent poses a new question for the first time, where to "pose" a question Q is to acquire of a view about Q. The mapping $\mathbf{I} \mapsto \mathbf{I} + Q^Q$ is the natural formalisation of this idea: its output $\mathbf{I} + Q^Q$ is the smallest extension of the information state \mathbf{I} that includes a view about Q.

This idea of deductive inquiry as a question-guided endeavour has a precedent in the slave boy from Plato's *Meno* (compare also Pérez Carballo 2016, Friedman 2017, 2020). Guided by the questions that

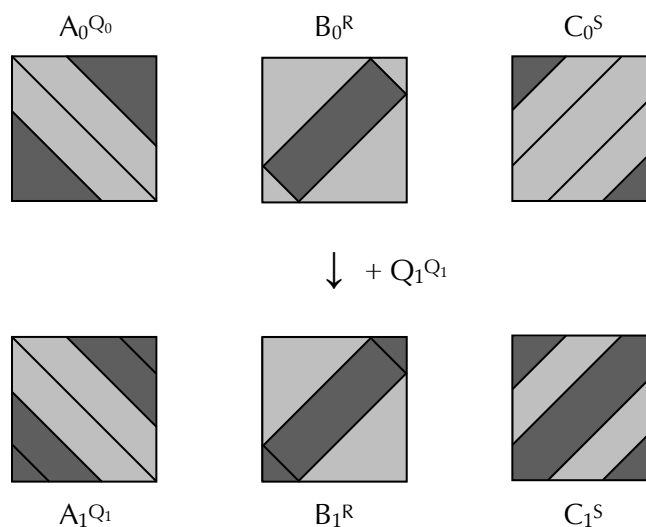


FIGURE 6: A TAUTOLOGICAL UPDATE ON OVERLAPPING VIEWS

Socrates asks him, the boy reasons his way to the conclusion that the diagonal of a square of size one is equal to the side of a square of size two. At the outset, the boy already has all the basic geometric intuitions he needs to figure this out. But he only arrives at the right conclusion after thinking through Socrates' strategically posed questions.

There are limitations to this paradigm. Tautological updates are a neat model for simple deductive inferences, but not every deduction can be modelled using tautological updates alone. For instance, the hypothetical reasoning in a proof by cases or in a *reductio ad absurdum* appears to involve the tentative addition of a piece of information to the subject's stock of beliefs. Moreover, as discussed in Section III, deductive inference can bring to light an inconsistency in the reasoner's beliefs, and force them to discard one of their prior beliefs. In cases like that, the deductive process involves belief revisions as well as updates.

VII. Three Kinds of Hard Questions

As the story of Socrates and the slave boy illustrates, human beings can increase their knowledge by posing new questions, rather than acquiring new information. But that ability is not unlimited. In this section, I discuss three different kinds of limits on our question-posing ability, each of which captures a different cognitive barrier curbing our deductive abilities. Bounds on our ability to pose *novel* questions capture our *conceptual* limitations. Bounds on our ability to pose very *large* questions capture our *computational* limitations. And finally, bounds on our ability to identify *good* questions capture our *strategic* limitations.

Novel Questions

In distinguishing between possibilities, we draw on our conceptual resources and world knowledge. Limitations in those resources limit what questions we can pose. One illustration of this point has come up several times already: King William believes he can avoid war with France. And as we saw in the last section, this means he is only a tautological update away from believing that he can avoid nuclear war with France. However, William is in no position to form a view on the requisite question:

Will England have a nuclear war with France or some other kind of war or no war at all. The reason is that William lacks the requisite conceptual resources: he does not know what a nuclear war is. (Seth Yalcin has a more detailed discussion of this point in Yalcin 2011, §6-8.)

Concepts might also facilitate the posing of new questions in cases where the subject does already have the ability to distinguish the relevant possibilities. One possible example is suggested in Pérez Carballo 2016. The Königsberg Bridges problem asks whether it is possible to take a roundtrip through the city of Königsberg that crosses each of its seven bridges exactly once. In Pérez Carballo's telling, Euler solved this problem by first posing a new question: *What is the graph-theoretic representation of the bridges and landmasses of Königsberg?* Each cell of this question is just an intensional proposition about the layout in the city, so that each one of these possibilities could also be described by someone who lacked the concept of a graph. Still, having the concept of a graph certainly makes it a great deal easier to partition the possibilities in this particular way, and it may well be essential.

Big Questions

I have a terrible sense of direction. When I am new in a city, I will figure out how to get from the hotel to the central square, say, and how I can get from the central square to the museum or the river bank, and from the museum to the restaurant. But having gathered all that information, I still will not be able to work out a half-way efficient route back from the restaurant to the hotel. Without a map, the safe option is just to retrace my steps: return to the museum, then back to the central square, then to the hotel. Else I will probably get lost. My friend is different. Given the exact same information, she will identify the shortest way back in a heartbeat, even if it runs through a neighbourhood she has never seen before.

Maybe my friend has a better memory for these things than I do. But what is more important is the way she puts all the information together. My web of beliefs about the city's geography is a chaotic patchwork of partially overlapping little maps, patched up with landmarks, mnemonics and other crutches. I only have answers to small, local questions about the city's geography, none of which have any bearing on unexplored areas. My friend, on the other hand, sees the bigger picture. Her views are more robustly connected because she has global views about the overall layout of the city that integrate her detailed views about the small parts. This puts her in a position to see we have walked in a big circle, and that the hotel is just a few blocks away. Her geographical beliefs answer bigger questions and are better linked together.

Maintaining a high level of integration between geographical beliefs is a non-trivial cognitive skill: my friend is better at it than I am. The skill can be improved with practice. London cab drivers are an extreme example. Over three to four years of intensive training for a harrowing exam called *The Knowledge*, they acquire the ability to efficiently deploy a vast amount of detailed geographical information in order to determine the fastest route from one place in London to another. It has been shown that in successful trainees, this learning process results in a significantly enlarged posterior hippocampus (Maguire et al. 2000, 2011). This is a remarkable illustration of the way that neuroplasticity allows human beings to go beyond their innate cognitive endowments. At the same time, the discovery that this requires additional grey matter implies that there is a limit to how far our abilities can be stretched. There is only so much new grey matter one can acquire, if only because there is a finite amount of space in a human skull.

Minimally rational agents see more consequences of their beliefs as they bring their views to bear on bigger questions, and as their views become better connected to other questions in the web. The more an agent's views are connected, the better information is distributed across the web, and the easier it is to recall. But there is a limit to the number of questions a finite agent can have views about, and also to the size of the questions. This limitation puts ideal rationality out of reach for us mere mortals: a minimally rational agent who had views on every question would be ideally rational, in that their beliefs would satisfy *Ideal Closure*, *Ideal Adjunction* and *Ideal Consistency*.¹⁴

Tautological updates take effort because they enlarge a believer's web of questions, and require the integration of various beliefs on different questions. Sometimes, a small question can already produce a sweeping cognitive change. But adding bigger questions with more parts is always more demanding: bigger questions contain the smaller ones as parts, making more distinctions and forging more connections. We have been thinking of parts as "free" consequences, which are believed without additional cogitation. But that is just to say that no cogitation *beyond* the acquisition of a belief in the whole is needed to believe the parts. The flip-side is that acquiring beliefs with many parts is hard, because it is a precondition for doing so that you acquire beliefs about the parts as well.

Good Questions

A good question can be hard to find. Since any tautological update requires cognitive effort to process, and resources are limited, we cannot look into *every* possible question. Consequently, when we engage in deductive reasoning, we inevitably make choices about which questions to look into. Sometimes, the hardest part of a deduction is not the update itself, but knowing which update to perform. Hitting on the right question may require insight or luck. In the *Meno*, Socrates' strategic questioning helps the slave boy precisely because it relieves him of the most creative part of the deductive process.

Here is an example taken from the psychology literature to illustrate the point (Levesque 1986, Toplak and Stanovich 2002). Based on the following three pieces of information, can you say whether or not an unmarried person is looking at a married person?

- 9) Jack is looking at Kate and Kate is looking at George
- 10) Jack is unmarried
- 11) George is married

Take a moment to picture the situation and think it through.

In Toplak and Stanovich's survey, 86% of subjects answered that the correct answer cannot be determined on the basis of the information provided. But as a matter of the fact it can. This becomes easy to see once you are given the following hint:

- 12) Either Kate is married or she is unmarried

Once those two possibilities are separated, the answer becomes clear. If Kate is married, then Jack is an

¹⁴ *Proof.* Suppose Laplace is a minimally rational agent with views on every question. Then Laplace's beliefs satisfy *Ideal Closure*. For suppose she believes A^Q , and suppose A^Q entails B^R . Since she has a view on QR , harmony demands she must believe AR^{QR} , which contains B^R as a part. So she believes B^R . Laplace's beliefs also satisfy *Ideal Adjunction*. For suppose she believes A^Q and B^R . She has a view on QR , so by harmony she believes AR^{QR} , QB^{QR} , and therefore AB^{QR} . Finally, Laplace's beliefs satisfy *Ideal Consistency*. If not, by *Ideal Adjunction*, she would believe \perp^Q for some Q and thus, by harmony, believe $\perp^{[Q]}$.

unmarried person looking at a married person, because Jack is looking at Kate. If Kate is unmarried, then she herself is an unmarried person looking at a married person, because she is looking at George. It is striking how an instance of the law of the excluded middle transforms an otherwise elusive inference into a no-brainer. This makes (12) an unusually simple and elegant example of an informative tautology.

We can account for this as follows. Conjoining the the three given premises (9-11) is insufficient to arrive at the conclusion that an unmarried person is looking at a married person. But this conclusion *is* a direct consequence of the conjunction of (9-12). To see this, associate premises (9), (10) and (11) with the quizpositions A^L , B^I and C^G respectively where

L: *Out of Jack, Kate and George, who is looking at whom?*

J: *Is Jack married?*

G: *Is George married?*

The task confronts subjects with something like the following question:

Q: *Who out of Jack, Kate and George is unmarried? And who is looking at a married person?*

And the target conclusion is:

D^Q: *One of Jack, Kate or George is an unmarried person looking at a married person.*

In approaching this problem, the natural first step is to try to picture the situation, putting all the given information together into a single representation. This can be modelled as an update with the tautologous quizposition LJG^{LJG} , which takes a state in which A^L , B^I and C^G are believed individually to a state where their conjunction ABC^{LJG} is also believed. However, doing this is not sufficient. The conjunction ABC^{LJG} entails the conclusion D^Q. But because Q is not part of LJG, it does not *contain* D^Q as a part. A further step is required to get from the belief ABC^{LJG} to the conclusion D^Q.

Other attempts to get at the answer also fail to work in this case. For instance, it is natural to break up Q into simpler questions: *Is Jack an unmarried person looking at a married person, Is Kate an unmarried person looking at a married person and Is George an unmarried person looking at a married person.* ABC^{LJG} fails to settle the first two questions, and entails a negative answer to the latter. These discouraging results would reasonably lead one to conclude that the given information is insufficient to settle Q. Plausibly, that is where the inquiry ends for most of Toplak and Stanovich's respondents.

We can only get at the target conclusion by making a further distinction. We need to separate two possibilities that LJG joins together: namely the possibility that *Jack and Kate are unmarried and only George is married*, and the possibility that *only Jack is unmarried and Kate and George are married*. Separating these two scenarios involves conjoining ABC^{LJG} with the content of (12), *Either Kate is married or not*. This is the tautologous quizposition K^K , where

K: *Is Kate married?*

After that further adjunction, the subject's overall view of the situation $ABCK^{LJGK}$; and since $LJGK$ *does* contain Q as a part, it follows that they now believe the target conclusion D^Q as well.

There is some amount of cognitive effort involved in making the extra distinction that takes you from ABC^{LJG} to $ABCK^{LJGK}$. But the students Toplak and Stanovich interviewed could all have made this

further reasoning step if prompted. The explanation for why they mostly failed to do so does not lie in the intrinsic difficulty of this update. Rather, the students must have *overlooked* the question K for some reason. It may be that it did not occur to them: there is experimental evidence that reasoners are better at *recognising* a good question when it is presented to them than they are at coming up with good questions on their own (Rothe, Lake and Gureckis 2018). It is also possible that they did consider the question, but made an *a priori*, metacognitive judgment that it was not worth looking into.

In this particular context, two factors may contribute to that decision. Firstly, the fact that they were not given information about Kate's marital state could be taken as an indication that this issue is irrelevant. There is independent evidence that reasoners are generally reluctant to think through a question when they know in advance that the answer is unknown (Tversky and Shafir 1992, Shafir 1994). Secondly, the fact that other lines of inquiry do not resolve the matter may have given rise to an overriding impression that further cogitation was pointless.

Because we cannot ask every question, we inevitably need heuristics to decide which questions to ignore and which to consider, and those must be prior to actually thinking through the questions. For instance, conjoining $ABC \cup G$ with the tautology *Either George owns a sloop or he does not* also make some new consequences available. But since you know *a priori* that those consequences will be of no help in resolving the task at hand, you would never look into that question: it is safely ignored. In the case at hand, the question K apparently does not appear fruitful to most people, although that appearance is misleading.

No doubt, these last two sections have raised as many questions about deductive inference as they answered. How do we model belief revision in this context? Can we give an informative quantitative measure of the difficulty of a given tautological update? How do reasoners assess the interest of a question prior to updating? But I hope I have said enough to show that the present account of minimally rational belief states gives us a systematic framework in which these and other theoretical questions about deduction can be fruitfully discussed, and relative to which experimental results about deductive inference can be sensibly interpreted. And because the old, intensional account of belief contents was ill-adapted to those tasks, this development opens up a whole new domain of inquiry for unstructured conceptions of belief content.

References

- Berto, Francesco, 2019, "Simple Hyperintensional Belief Revision." *Erkenntnis* 84 (3):559-75.
- Borgoni, Cristina, Dirk Kindermann, and Andrea Onofri (eds.), 2021, *The Fragmented Mind*. Oxford: Oxford University Press.
- Breen, Mara, 2014, "Empirical Investigations of the Role of Implicit Prosody in Sentence Processing." *Language and Linguistics Compass* 8(2): 37-50.
- Bjerring, Jens Christian, and Mattias Skipper, 2019, "A Dynamic Solution to the Problem of Logical Omniscience." *Journal of Philosophical Logic* 48: 501-21.
- Cherniak, Christopher, 1986, *Minimal Rationality*. Cambridge, MA: Bradford Books.
- Ciardelli, Ivano and Floris Roelofsen, 2015, "Inquisitive dynamic epistemic logic." *Synthese*, 192(6), 1643-87.
- Ciardelli, Ivano, Jeroen Groenendijk, and Floris Roelofsen, 2019, *Inquisitive Semantics*. Oxford University Press.
- Cresswell, Max, 1975, "Hyperintensional logic." *Studia Logica* 34(1): 25 - 38.
- 1985, *Structured Meanings*. Cambridge, MA: MIT Press.
- 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.

- Eberle, Rolf A., 1974, "A logic of believing, knowing, and inferring." *Synthese* 26 (3-4): 356-382.
- Egan, Andy, 2008, "Seeing and believing: perception, belief formation and the divided mind." *Philosophical Studies* 140(1): 47-63
- Elga, Adam, and Agustín Rayo, 2021. "Fragmentation and information access." In Borgoni, Kindermann and Onofri (eds.) 2021.
2022, "Fragmentation and logical omniscience." *Noûs* 56(3): 716-41.
- Egré, Paul and Denis Bonnay, 2012, "Metacognitive perspectives on unawareness and uncertainty." In: *Foundations of Metacognition*, eds. M.J. Beran et al., New York: Oxford University Press.
- Fagin, Ronald and Joseph Y. Halpern, 1988, "Belief, awareness, and limited reasoning." *Artificial Intelligence* 34: 39-76
- Fagin, Ronald, Joseph Y. Halpern and Moshe Y. Vardi, 1995, "A Nonstandard Approach to the Logical Omniscience Problem." *Artificial Intelligence* 79(2): 203-240.
- Fine, Kit, 2013, "A note on partial content." *Analysis* 73(3): 413-419.
2016, "Angelic Content." *Journal of Philosophical Logic* 45(2): 199-226.
2017, "A Theory of Truthmaker Content I, II" *J. of Philosophical Logic* 46(6): 625-702.
- Friedman, Jane, 2020, "The Epistemic and the Zetetic." *The Philosophical Review* 129(4): 501-536.
2017, "Inquiry and Belief." *Noûs* 53(2): 296-315.
- Fritz, Peter, and Harvey Lederman, 2015, "Standard State-Space Models of Unawareness." *Proceedings of TARK XV: Fifteenth Conference on Theoretical Aspects of Rationality and Knowledge*, Pittsburgh.
- Goodman, Jeremy, 2019, "Agglomerative Algebras." *Journal of Philosophical Logic* 48: 631-48.
- Greco, Daniel, 2015, "Iteration and Fragmentation." *Phil. and Phenomenological Research* 91(3): 656-673.
- Groenendijk, Jeroen and Martin Stokkhhof, 1997, "Questions." In *Handbook of Logic and Language*, Johan van Benthem and Alice ter Meulen (eds.), 1055-1124. Amsterdam: Elsevier 1997.
- Hamblin, C.L., 1958, "Questions." *Australasian Journal of Philosophy* 36, 159-168.
- Hawke, Peter, 2016, "Questions, topics and restricted closure." *Philosophical Studies* 173: 2759-2784
- Hintikka, Jaakko, 1992, *Knowledge and Belief*. Ithaca: Cornell University Press.
- Hoek, Daniel, 2019, *The Web of Questions: Inquisitive Decision Theory and The Bounds of Rationality*. PhD dissertation, New York University.
2022, "Questions in Action." *Journal of Philosophy*.
- Holguín, Ben, 2022, "Thinking, Guessing and Believing." *Philosophers' Imprint* 22(6): 1-25.
- Jago, Mark, 2013, "The Content of Deduction." *Journal of Philosophical Logic* 42(2): 317-334.
2014, *The Impossible*. Oxford: Oxford University Press.
2015, "Hyperintensional Propositions." *Synthese* 192 (3): 585-601.
- Konolige, Kurt, 1986, *A Deduction Model of Belief*. Burlington, MA: Morgan Kaufmann.
- Koralus, Philipp, and Salvador Mascarenhas, 2013, "The Erotetic Theory of Reasoning" *Philosophical Perspectives* 27(1): 312-365.
2018, "Illusory inferences in a question-based theory of reasoning." In: *Pragmatics, Truth, and Underspecification: Towards an Atlas of Meaning*, eds. Ken Turner and Laurence Horn, 300-322. Leiden: Brill.
- Levesque, Hector J., 1984, "A logic of implicit and explicit belief." *AAAI-84 Proceedings*, 198-202.
1986, "Making believers out of computers." *Artificial Intelligence* 30, 81-108.
- Lewis, David, 1988, "Relevant Implication." *Theoria* 54(3): 161-74.
- Maguire, E.A., D.G. Gadian, I.S. Johnsrude, C.D. Good, J. Ashburner, R.S. Frackowiak, C.D. Frith, 2000 "Navigation-related structural change in the hippocampi of taxi drivers." *Proceedings of the National Academy of Sciences USA* 97: 4398-4403.
- Maguire, E.A. and K. Woollett, 2011, "Acquiring 'the Knowledge' of London's Layout Drives Structural Brain Changes." *Current Biology* 21(24): 2109-2114.
- Montague, Richard, 1970, "Universal Grammar." *Theoria* 36(3): 373-398.
- Moore, Robert C. and Gary G. Hendrix, 1979, "Computational Models Of Beliefs and The Semantics Of Belief-Sentences." Technical Note 187, SRI International, Menlo Park.
- Norby, Aaron, 2014, "Against Fragmentation." *Thought* 3(1): 30-38.
- Pérez Carballo, Alejandro, 2016, "Structuring logical space." *Phil. and Phenomenological Research* 92(2): 460-491.
- Rothe, Anselm, Brenden M. Lake and Todd M. Gureckis 2018, "Do people ask good questions?" *Computational Brain & Behavior*, 1, 69-89.
- Rooth, Mats, 1992, "A theory of focus interpretation." *Natural Language Semantics* 1: 75-116.
- Schaffer, Jonathan, 2007, "Knowing the answer." *Philosophy and Phenomenological Research*, 75(2): 383-403.

- Schipper, Burkhard C., 2015, "Awareness." In: *Handbook of Epistemic Logic*, eds. Hans van Ditmarsch et al., 77-146. London: College Publications.
- Scott, Dana, "Advice on modal logic." In *Philosophical problems in logic*, ed. Karel Lambert, 143-173. Dordrecht: D. Reidel.
- Shafir, Eldar, 1994, "Uncertainty and the difficulty of thinking through disjunctions." *Cognition* 50: 403-430.
- Sim, Kwang Mong, 1997, "Epistemic logic and logical omniscience: A survey." *International Journal of Intelligent Systems* 12(1): 57-81.
- Solaki, Anthia, Francesco Berto and Sonja Smets, 2019, "The Logic of Fast and Slow Thinking." *Erkenntnis*.
- Stalnaker, Robert, 1984, *Inquiry*. Cambridge, MA: MIT Press.
- 1999, "The Problem of Logical Omniscience II." In his *Context and Content*, 255-273. New York: Oxford University Press.
- Toplak, Maggie and Keith Stanovich, 2002, "The Domain Specificity and Generality of Disjunctive Reasoning." *Journal of Educational Psychology* 94(1): 197-209.
- Tversky, Amos and Eldar Shafir, 1992. "The disjunction effect in choice under uncertainty." *Psychological Science* 3: 305-309.
- Van Fraassen, Bas, 1969, "Facts and Tautological Entailments." *Journal of Philosophy* 66(15): 477-487.
- Vardi, Moshe Y., 1986, "On Epistemic Logic and Logical Omniscience." In: *Proc. of TARK 1986*, Joseph Y. Halpern (ed.), 293-306. Los Altos, CA: Morgan Kaufmann.
- Yablo, Stephen, 2014, *Aboutness*. Princeton, NJ: Princeton University Press.
- 2017, "Open knowledge and changing the subject." *Philosophical Studies* 174(4): 1047-71.
- Yalcin, Seth, 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 Borgoni, Kindermann and Onofri (eds.) 2021.

Appendix: Belief Reports

Section IV above articulated and defended three general principles of minimally rational belief: *Closure under Parthood*, *Partial Adjunction* and *Partial Consistency*. My official interpretation of these principles is as schemata that range over English sentences, using the following abbreviations:

\wedge, \vee, \neg	and, or, not
$B\phi$	α believes that ϕ
$\phi \models \psi$	ϕ semantically entails ψ
$\phi \leq \psi$	the quizposition ϕ expresses contains the quizposition ψ expresses
$\phi \approx \psi$	the quizpositions ϕ and ψ express are about the same question [or alternatively: $\phi \leq (\psi \vee \neg\psi)$ and $\psi \leq (\phi \vee \neg\phi)$]

For instance, *Closure under Parthood* says, in unabbreviated form, that if the quizposition expressed by ϕ contains the quizposition expressed by ψ , then $\lceil \alpha$ believes that $\phi \rceil$ entails $\lceil \alpha$ believes that $\psi \rceil$. This appendix provides a semantics of belief reports that validates these schemata, while also invalidating *Ideal Closure*, *Ideal Adjunction* and *Ideal Consistency*.

For the reasons explained in Section I, this is impossible if the only meaning we assign the preajcent of a belief report is its truth conditions. So we shall assume declarative sentences are also associated with a compositionally determined question. There is independent linguistic motivation for such an approach, and there are in fact numerous well-developed semantic frameworks on the market that do something like this, notably alternatives semantics (Hamblin 1973, Rooth 1992), truthmaker semantics (Van Fraassen 1969, Yablo 2014, Fine 2017), and inquisitive semantics (Ciardelli et al. 2019).

Let us write $[\phi]$ for the question assigned to ϕ and $\llbracket\phi\rrbracket$ for the quizposition expressed by ϕ (the question-answer pair). The semantic clauses for negation, conjunction and disjunction can be given in terms of the corresponding quizpositional operations, as defined in Section III:

$$\begin{aligned} \llbracket \text{not } \phi \rrbracket &= \neg \llbracket \phi \rrbracket \\ \llbracket \phi \text{ and } \psi \rrbracket &= \llbracket \phi \rrbracket \wedge \llbracket \psi \rrbracket \\ \llbracket \phi \text{ or } \psi \rrbracket &= \llbracket \phi \rrbracket \vee \llbracket \psi \rrbracket = \neg (\neg \llbracket \phi \rrbracket \wedge \neg \llbracket \psi \rrbracket) \end{aligned}$$

What about atomic sentences π ? The simplest course is to let $[\pi]$ be the polar question $\{p, \neg p\}$, where p is the set of worlds at which π is true. But we can do better if we take a suggestion from Yalcin 2011, and assume that the question component incorporates the focus alternatives of the sentence: if the set of focus alternatives of π is any set of intensional propositions A , define $[\pi]$ as the coarsest partition question such that each proposition in A is a union of $[\pi]$ -cells. For our purposes, the advantage of this approach is that it lets us capture parthood relations between atomic sentences. For instance, the quizposition expressed by “Jill’s birthday is in **April**” comes out as part of that expressed by “Jill’s birthday is on **April 15th**”. Sentences often have a contextually inferred focus even when they lack any audible or typographical stress (Breen 2014). But if an atomic sentence π lacks focus alternatives altogether, we can revert to the default, polar option for $[\pi]$.

Now for the semantics of “believe”. Since we are giving a quizpositional semantics, we will need to specify what question is answered by a sentence of the form “Jake believes that ϕ ”. I suggest the natural candidate is the question *What are Jake’s beliefs about Q?*, where Q is the question the prejacent ϕ is about. Let’s do this in steps. First, to zero in on the relevant set of beliefs, let us define:

$$\mathbf{DOX}(X, Q, w) = \{A^R \in \mathbf{B}_{X, w} : R \text{ is part of } Q\}$$

where $\mathbf{B}_{X, w}$ is the inquisitive belief state of agent X at world w . Next, let $\Theta(X, Q)$ be the partition induced by the following equivalence relation on worlds: $w \sim v$ iff $\mathbf{DOX}(X, Q, w) = \mathbf{DOX}(X, Q, v)$. Finally, we specify the clause for “believe” as follows (where “ α ” names X):

$$\begin{aligned} [\alpha \text{ believes that } \phi] &= \Theta(X, [\phi]) \\ \llbracket \alpha \text{ believes that } \phi \rrbracket &= \langle \Theta(X, [\phi]), \{t \in \Theta(X, [\phi]) : \forall w \in t, \llbracket \phi \rrbracket \in \mathbf{DOX}(X, [\phi], w)\} \rangle \end{aligned}$$

This semantics validates *Closure under Parthood*, *Partial Adjunction* and *Partial Consistency*, and hence also *Closure under Conjunction Elimination*, *Avoidance of Contradictions* and *Avoidance of Contradictories*. At the same time, *Ideal Closure*, *Intensionality*, *Ideal Adjunction* and *Ideal Consistency* are all invalidated.

Other notable validities under this semantics include the following:

$$\begin{aligned} &\textit{Closure under Material Modus Ponens. } B\phi, B(\neg\phi \vee \psi) \models B\psi \\ &\textit{Strong Avoidance of Contradictories. } \text{If } \phi \not\models \psi, \text{ then } B\phi \models \neg B\neg\psi \\ &\textit{Closure under Embedded Parthood. } \text{If } \phi \leq \psi, \text{ then } BB\phi \models BB\psi \end{aligned}$$

Closure under Material Modus Ponens and *Strong Avoidance of Contradictories* further heighten the contrast of the present account with fragmentation theories of belief, which lack non-trivial coherence principles of this sort.