Belief in Truthmaker Semantics

Mark Jago

Abstract: Stephen Yablo has argued that ascriptions of belief and knowledge are sensitive to *subject matter* and that theorising in such terms may help resolve philosophical questions about the semantics of such ascriptions. *Truthmaker semantics* offers a way of theorising about subject matters. My main aim in this paper is to investigate what a semantics for knowledge or belief ascriptions might look like within truthmaker semantics. I then discuss what the resulting account might have to say about the problem of logical omniscience, the nature of concepts, and rational belief.

1   Introduction

This paper has two main aims. The first is to investigate what a semantics for knowledge or belief ascriptions looks like within the setting of *truthmaker semantics*. The second is to evaluate how well that approach does as a solution to the various problems of *logical omniscience*, whereby agents are modelled as knowing too many logical consequences of what they know.

   I will focus on belief, rather than knowledge, but much of what I say here applies equally to knowledge ascriptions. More importantly, I will situate the discussion within *Finean* truthmaker semantics (Fine 2017a;b), as opposed to the *Yabloesque* approach (Yablo 2014). The former takes primitive *states*, standing in a part-whole relation, as its conceptual starting point. The latter, by contrast, takes possible worlds as its primitive and constructs states from these. The former is more general, in that it can draw hyperintensional distinctions where the latter cannot (Yablo 2018; Fine 2020). It is also more foreign. We know far less about how to reason with and about states than we do with and about worlds. That is reason enough to explore.

   The basic insight on how to model knowledge and belief in this setting, however, is from Yablo:

> Knowledge-attributions care about subject matter, over and above truth conditions. They take note of how *P* is true or false in various worlds, not only which world is it is true or false in. (Yablo 2014, 122)

You know you locked the front door. Do you thereby know that any apparent evidence to the contrary is misleading? It seems not. Reports of an open door might make you reconsider. So if we are not to concede too much everyday knowledge to the sceptic, we must accept that knowledge is not closed under simple known implications such as this. Yablo's diagnosis highlights the change in subject matter, from *how the door is* to (something like) *whether there is evidence from other sources concerning how the door is* (Yablo 2014, §7.4).

   The paper proceeds as follows. §2 is a brief introduction to truthmaker semantics, *à la* Fine. §3 presents the basic semantics for belief ascriptions, which

is then extended to a more general account in §4. Additional conditions – *truth* and *positive introspection* – are discussed in §5. We then turn, in §6, to how the resulting account handles the various problems of logical omniscience. §7 relates the account to some recent approaches to *concept possession*. §8 discusses an objection which affects similar responses to logical omniscience. Finally, §9 discusses the vexing issue of how an agent's beliefs stand in light of her limited cognitive resources.

2 Truthmaker Semantics

In this section, we sketch a very brief overview of truthmaker semantics and state some basic results. For further details, the reader is referred to Fine 2016; 2017a and Fine and Jago 2019; Forthcoming. Here we use the notation and presentation of Fine and Jago Forthcoming.

Truthmaker semantics is built around the notion of a *state space*: a set of states S with a part-whole structure $\sqsubseteq$ on it. This is a *complete partial order*: it is reflexive and transitive and has unique least and greatest elements, $\square, \blacksquare \in S$, such that $\square \sqsubseteq s \sqsubseteq \blacksquare$ for every state $s \in S$. Given this ordering, each subset T of U has a unique *least upper bound* $\bigsqcup T \in S$ and a unique *greatest lower bound* $\bigsqcap T \in S$. (The least upper bound of a set T is the least state $s \in S$ for which $t \sqsubseteq s$ for each $t \in T$. Similarly, the greatest lower bound of a set T is the greatest state $s \in S$ for which $s \sqsubseteq t$ for each $t \in T$.) In particular, $\bigsqcup \varnothing = \bigsqcap S = \square$ and $\bigsqcup S = \bigsqcap \varnothing = \blacksquare$. For pairs of states, we write $s \sqcup u$ for $\bigsqcup \{s, u\}$, the *fusion* of s and u, and $s \sqcap u$ for $\bigsqcap \{s, u\}$. $\langle S, \sqcup \rangle$ and $\langle S, \sqcap \rangle$ are *complete semilattices*, with identity elements $\square$ and $\blacksquare$, respectively: $s \sqcup \square = s$ and $s \sqcap \blacksquare = s$ for each $s \in S$. $\sqcup$ and $\sqcap$ are commutative, associative, and idempotent, and $s \sqcup u = u$ iff $s \sqcap u = s$ iff $s \sqsubseteq u$. (In what follows, we mostly ignore $\sqcap$.)

We shall work with a very simple propositional language:

$$p \quad | \quad \neg A \quad | \quad A \wedge B \quad | \quad A \vee B$$

(to be expanded with belief operators B_i in §3.) We expand state spaces $\langle S, \sqsubseteq \rangle$ to models by adding *positive* and *negative valuation* functions:

Definition 1 (Models)**.** A *model* $\mathcal{M}$ is a quadruple $\langle S, \sqsubseteq, V^+, V^- \rangle$, where $\langle S, \sqsubseteq \rangle$ is a state space and $V^+, V^- : \mathcal{P} \longrightarrow 2^S$ are functions from sentence letters to nonempty subsets of S.

We then define exact truthmaking ($\Vdash$) and falsitymaking ($\dashv\!\vdash$) relations as follows:

Definition 2 (Exact truthmaking and falsitymaking)**.** Given a model $\mathcal{M}$ (which we leave implicit), *exact truthmaking* $\Vdash$ and *exact falsitymaking* $\dashv\!\vdash$ relations are defined by double recursion as follows:

$$s \Vdash p \quad \text{iff} \quad s \in V^+ p$$
$$s \dashv\!\vdash p \quad \text{iff} \quad s \in V^- p$$

$$
\begin{aligned}
s \Vdash \neg A \quad &\text{iff} \quad s \dashv\vert A \\
s \dashv\vert \neg A \quad &\text{iff} \quad s \Vdash A \\
s \Vdash A \wedge B \quad &\text{iff} \quad \exists tu(s = t \sqcup u \ \& \ t \Vdash A \ \& \ u \Vdash B) \\
s \dashv\vert A \wedge B \quad &\text{iff} \quad s \dashv\vert A \text{ or } s \dashv\vert B \\
s \Vdash A \vee B \quad &\text{iff} \quad s \Vdash A \text{ or } s \Vdash B \\
s \dashv\vert A \vee B \quad &\text{iff} \quad \exists tu(s = t \sqcup u \ \& \ t \dashv\vert A \ \& \ u \dashv\vert B)
\end{aligned}
$$

These are said to be *exact* relationships in that $s \Vdash A$ says that s is wholly relevant to A's truth (and $s \dashv\vert A$ says that s is wholly relevant to A's falsity). This requirement leads to the unusual truthmaking clause for conjunction: a truthmaker for $A \wedge B$ is the fusion of a truthmaker for A and a truthmaker for B. That state may itself not be *wholly* relevant to A's (or B's) truth and so, in general, truthmakers for conjunctions will not be truthmakers for their conjuncts. (Similar remarks apply to falsitymaking for disjunctions.)

That state will nevertheless be *sufficient* for (if not wholly relevant to) the truth of the conjuncts. This is the notion of *inexact truthmaking*, defined as follows.

Definition 3 (Inexact truthmaking and falsitymaking). In any model $\mathcal{M}$, $s \Vvdash A$ iff $u \Vdash A$ for some $u \sqsubseteq s$, and $s \mathrel{\vert\!\!\vartriangleleft} A$ iff $u \dashv\vert A$ for some $u \sqsubseteq s$.

Inexact truthmaking, unlike its exact cousin, obeys the standard extensional clauses for conjunction and disjunction:

$$
\begin{aligned}
s \Vvdash A \wedge B \quad &\text{iff} \quad s \Vvdash A \text{ and } s \Vvdash B \\
s \mathrel{\vert\!\!\vartriangleleft} A \wedge B \quad &\text{iff} \quad s \mathrel{\vert\!\!\vartriangleleft} A \text{ or } s \mathrel{\vert\!\!\vartriangleleft} B \\
s \Vvdash A \vee B \quad &\text{iff} \quad s \Vvdash A \text{ or } s \Vvdash B \\
s \mathrel{\vert\!\!\vartriangleleft} A \vee B \quad &\text{iff} \quad s \mathrel{\vert\!\!\vartriangleleft} A \text{ and } s \mathrel{\vert\!\!\vartriangleleft} B
\end{aligned}
$$

We define *exact truthmaker* and *falsitymaker sets* as follows:

$$
|A|^{+} = \{s \in S \mid s \Vdash A\} \qquad |A|^{-} = \{s \in S \mid s \dashv\vert A\}
$$

and lift $\sqcup$ to sets of states by setting:

$$
T \sqcup U = \{t \sqcup u \mid s \in T, u \in U\}
$$

We may then state the exact clauses in algebraic form:

$$
\begin{aligned}
|\neg A|^{+} = |A|^{-} \qquad &|A \wedge B|^{+} = |A|^{+} \sqcup |B|^{+} \qquad |A \vee B|^{+} = |A|^{+} \cup |B|^{+} \\
|\neg A|^{-} = |A|^{+} \qquad &|A \wedge B|^{-} = |A|^{-} \cup |B|^{-} \qquad |A \vee B|^{-} = |A|^{-} \sqcup |B|^{-}
\end{aligned}
$$

We identify *propositions* with sets of states. More precisely, a *unilateral proposition* is a set of states $P \subseteq S$, and a *bilateral proposition* is a pair $\mathbf{P} = \langle P^{+}, P^{-} \rangle$ of sets of states. The idea here is that P^{+} contains $\mathbf{P}$'s truthmakers and P^{-} its

falsitymakers. Given bilateral propositions $\mathbf{P} = \langle P^+, P^- \rangle$ and $\mathbf{Q} = \langle Q^+, Q^- \rangle$, we define bilateral Boolean operators as follows:

$$\neg \langle P^+, P^- \rangle = \langle P^-, P^+ \rangle$$
$$\langle P^+, P^- \rangle \wedge \langle P^+, P^- \rangle = \langle P^+ \sqcup Q^+, P^- \cup Q^- \rangle$$
$$\langle P^+, P^- \rangle \vee \langle P^+, P^- \rangle = \langle P^+ \cup Q^+, P^- \sqcup Q^- \rangle$$

We may in addition require one or more closure conditions on propositions:

Definition 4 (Closure Conditions).

CLOSURE ($\sqcup$): A proposition P is *closed* when $\bigsqcup Q \in P$ for any nonempty $Q \subseteq P$

CONVEX CLOSURE ($\frown$): A proposition P is *convex* when, for any $t \in S$, if $s, u \in P$ and $s \sqsubseteq t \sqsubseteq u$, then $t \in S$ too.

REGULAR CLOSURE ($*$): A proposition P is *regular* when it is both closed and convex.

We write $P^\sqcup$, $P^\frown$, and P^* for the smallest closed, convex, and regular sets (respectively) that contain P.

Each proposition has a *subject matter*, which intuitively is what the proposition is *about*. For Yablo (2014; 2018), each instance of 'there are n stars' has the subject matter *the number of stars*, analogous to the question, 'how many stars are there?' Fine (2017b) has a slightly different notion of subject matter, on which each instance of 'there are n stars' has the subject matter *whether there are n stars*, which (for the same n) it shares with 'there are not n stars'. Both agree that a proposition's subject matter is not given merely by the *objects* it is about, for the headlines, 'man bites dog' and 'dog bites man' have quite different subject matter (Yablo 2014, 24).

Following Fine (2017b), we define the subject matter $\mathbf{p}$ of a unilateral proposition P to be $\bigsqcup P$ and of a bilateral proposition $\mathbf{P} = \langle P^+, P^- \rangle$ to be $\mathbf{p}^+ \sqcup \mathbf{p}^-$. (We could instead understand bilateral subject matter as the *pair* $\langle \mathbf{p}^+, \mathbf{p}^- \rangle$. Fine (2017b, 697) calls these options 'comprehensive' and 'differentiated' subject maters, respectively. We adopt the former here because it captures the intuitive principle that negating a proposition does not affect its subject matter.)

Given a model $\mathcal{M}$ and sentence A, we shall for the most part be interested in the *regular* (unilateral and bilateral) propositions $|A|^{+*}$ and $\langle |A|^{+*}, |A|^{-*} \rangle$ associated with A. We use the notation $sm^+(A)$ and $sm^\pm(A)$ for their subject matters, respectively. These are determined purely by the subject matters of the letters appearing in A (and, for the unilateral case, by whether those letters occur *positively* or *negatively*, in the following sense).

Lemma 2.1 (Subject matter). *Say that p occurs* positively *(or* negatively*) in A when p occurs within the scope of an even (odd) number of negations in A. Let $lett^+(A)$ and $lett^-(A)$ be the sets of letters occurring positively and negatively in A, respectively, and set $lett(A) = lett^+(A) \cup lett^-(A)$. Then, relative to any model $\mathcal{M}$:*

4

(i) $\quad sm^{\pm}(A) = \bigsqcup_{p \in lett(A)} sm^{\pm}(p)$

(ii) $\quad sm^{+}(A) = \bigsqcup_{p \in lett^{+}(A)} sm^{+}(p) \ \sqcup \bigsqcup_{p \in lett^{-}(A)} sm^{+}(\neg p)$

Proof: (i) is by induction on A. The base case is given by definition and, for the induction step, it suffices to note that, for $\neg$: $sm^{\pm}(\neg A) = sm^{\pm}(A)$ and $lett(\neg A) = lett(A)$, and for $\wedge$ and $\vee$: $sm^{\pm}(A \wedge B) = sm^{\pm}(A \vee B) = sm^{\pm}(A) \sqcup sm^{\pm}(B)$ and $lett(A \wedge B) = lett(A \vee B) = lett(A) \cup lett(B)$.

For (ii), let $lit(A)$ be the set of *literals* (letters or their negations) occurring in A and $dnf(A)$ be any *disjunctive normal form* of A. DNFs have the property that $p \in lit(dnf(A))$ iff $p \in lett^{+}(A)$ and $\neg p \in lit(dnf(A))$ iff $p \in lett^{-}(A)$. Moreover, given the equivalences in figure 1, $dnf(A)$ is equivalent to A, hence $|dnf(A)|^{+*} = |A|^{+*}$, and so $sm^{+}(A) = sm^{+}(dnf(A))$. We also have

$$sm^{+}(A \wedge B) = sm^{+}(A \vee B) = sm^{+}(A) \sqcup sm^{+}(B)$$

and so:

$$sm^{+}(dnf(A)) = \bigsqcup_{l \in lit(dnf(A))} sm^{+}(l)$$
$$= \bigsqcup_{p \in lett^{+}(A)} sm^{+}(p) \ \sqcup \bigsqcup_{p \in lett^{-}(A)} sm^{+}(\neg p)$$

∎

We take sentences A, B to be *equivalent*, $A \equiv_{e} B$ when they express the same unilateral proposition. Where we impose no closure conditions on propositions, equivalence amounts to A, B having the same truthmakers. But for technical reasons, it is preferable to insist that propositions be regular closed sets, so that $A \equiv_{e} B$ when $|A|^{+*} = |B|^{+*}$. This gives us the familiar equivalences shown in figure 1. (Of these, the majority hold on the basic semantics. Idempotence for $\wedge$ requires closure. Distributivity for $\vee$ requires both closure and convexity.)

We take (single-premise) entailment to be propositional inclusion: A exactly entails B when $|B|^{+*} \subseteq |A|^{+*}$. As already noted, $A \wedge B$ will not exactly entail A. Yet there is clearly an important relationship between $A \wedge B$ and A. We say that the proposition expressed by A is a *conjunctive part* of that expressed by $A \wedge B$:

Definition 5 (Conjunctive parthood). P is a *conjunctive part* of Q, $P \leq Q$, when:

(Up) Each $s \in P$ is part of some $u \in Q$ (i.e. $s \sqsubseteq u$); and

(Down) Each $s \in Q$ has a part $u \in P$ (i.e. $u \sqsubseteq s$).

When the first condition holds, we say that P *subserves* Q ($P \sqsubseteq_{\forall \exists} Q$) and when the second is met, we say that Q *subsumes* P ($Q \sqsupseteq_{\forall \exists} P$). We also say that Q *contains* P when $P \leq Q$.

Note that $\leq$ is a natural way to lift the parthood ordering $\sqsubseteq$ from states to sets of states, since we have, in parallel to the usual order-lattice equivalence on states ($s \sqsubseteq u$ iff $s \sqcup u = u$):

Lemma 2.2 ([Fine 2017a](#)). $P \leq Q$ iff $P \wedge Q = Q$

Commutativity:	$A \wedge B \;\equiv_e\; B \wedge A$
	$A \vee B \;\equiv_e\; B \vee A$
Associativity:	$A \wedge (B \wedge C) \;\equiv_e\; (A \wedge B) \wedge C$
	$A \vee (B \vee C) \;\equiv_e\; (A \vee B) \vee C$
Distributivity:	$A \wedge (B \vee C) \;\equiv_e\; (A \wedge B) \vee (A \wedge C)$
	$A \vee (B \wedge C) \;\equiv_e\; (A \vee B) \wedge (A \vee C)$
Idempotence:	$A \wedge A \;\equiv_e\; A$
	$A \vee A \;\equiv_e\; A$
De Morgan:	$\neg(A \wedge B) \;\equiv_e\; \neg A \vee \neg B$
	$\neg(A \vee B) \;\equiv_e\; \neg A \wedge \neg B$
Double negation:	$\neg\neg A \;\equiv_e\; A$

Figure 1: Equivalences given regular closure

3 Semantics for Belief States

It is common in formal epistemology to take an agent's *total belief state* to be
a proposition (a set of worlds, situations, scenarios, states, or whatever) and
to analyse belief in a particular proposition, P, in terms of P's inclusion in (the
consequences of) that total belief state. We may adopt that approach in the
truthmaker setting, where we already have a notion of *proposition*. It remains to
say only what the relevant sense of *inclusion* is. Spoiler: it is *conjunctive parthood*.

We will begin with a simplified semantics, in this section, before extending to
the full analysis in §4. For the time being, we consider a single agent, for whom we
introduce a belief operator B into the language, so that BA is a sentence whenever
A is. An agent's total belief state is modelled as a unilateral proposition: a set
of states, D. The agent believes that A when the proposition expressed by A is
contained in (i.e. is a conjunctive part of) D. (This, I take it, is how to understand
Yablo's idea that '[belief]-attributions care about subject matter' (2014, 122)
within Finean truthmaker semantics.) More precisely:

Definition 6 (Simple models and truth)**.** A *simple doxastic model* is a quintuple
$\mathcal{M} = \langle S, \sqsubseteq, D, V^+, V^- \rangle$, where $\langle S, \sqsubseteq, V^+, V^- \rangle$ is as before and $D \subseteq S$ is a regular
unilateral proposition. For a B-free sentence A, BA is *true in* $\mathcal{M}$, $\mathcal{M} \vDash$ BA, when
$|A|^{+*} \leq D$. Entailment is preservation of truth-in-a-model and equivalence is
two-way entailment.

This definition of truth is very limited, applying only to *simple* belief ascriptions,
of the form BA where A itself is B-free. This is sufficient for a modest investigation
into the entailment behaviour of belief ascriptions, which shall be the focus of this
section. We shall give a more nuanced semantics in the next section.

As mentioned above, I take this approach to be the natural way of developing Yablo's suggestion within Finean truthmaker semantics. As far as I know, Fine does not suggest a semantics of belief or knowledge along these lines. Hawke and Özgün (2023) develop a truthmaker semantics for knowledge ascriptions along quite different lines: see Fine 2023 for discussion. Elgin (2021) connects knowledge ascriptions to a notion of *analytic consequence*, which may in turn be understood in terms of truthmaker semantics (see theorem 3.5 below), but does not offer a truthmaker semantics for knowledge ascriptions. So, as far as I am aware, this particular semantics for belief appears for the first time here. The closest existent approach I know of is Fine's analysis of *free-choice obligation* (Fine 2018), in which a deontic statement of obligation OA is true relative to a *code of conduct* C (a set of states) iff $OA \leq C$. The parallel to the present approach is not exact, however.

Under what conditions does BA entail BC, on this approach? The first source of closure conditions is the conditions on propositions themselves. If A and B express the same proposition, then BA will be equivalent to BC. So (having assumed the regular semantics) pairs BA, BC will be equivalent for each pair A, C listed in figure 1: $B(A \wedge C)$ is equivalent to $B(C \wedge A)$, and so on.

The second source of closure conditions is the relation of *conjunctive parthood* that TMS claims to hold between a believed proposition P and the agent's total belief state D. So let us look at some of the properties of conjunctive parthood in more detail.

Lemma 3.1 (Fine 2017a).

(i) $\leq$ *is reflexive and transitive*

(ii) *If* $P \leq Q$ *and* $Q \leq P$ *then* $P = Q$

(iii) *If* $P = P_1 \wedge P_2 \wedge \cdots$, *then (i)* $\mathbf{p} = \bigsqcup_i \mathbf{p}_i$ *and (ii)* P *is the least proposition such that* $P_i \leq P$ *for each* i

(iv) *If* $P \leq R$ *and* $Q \leq R$ *then* $P \wedge Q \leq R$

(v) $P \leq Q$ *iff* $\mathbf{p} \sqsubseteq \mathbf{q}$ *and* $Q \sqsupseteq_{\forall \exists} P$.

Part (v) will be especially useful in what follows, as it allows us to replace the (Down) condition on $P \leq Q$ with the simpler condition $\mathbf{p} \sqsubseteq \mathbf{q}$.

Given transitivity, if $|C|^{+*} \leq |A|^{+*}$, then BA will entail BC. In particular, $B(A \wedge C)$ entails both BA and BC. And conversely, by (iv), BA and BC together entail $B(A \wedge C)$. In fact, the agent's total belief state D may be seen as the big conjunction of all the propositions P_i she believes: $D = P_1 \wedge P_2 \wedge \cdots$. Its total subject matter is the fusion of the subject matters of each believed P_i. Thus, fixing D fixes a subject matter $\mathbf{d} = \bigsqcup D$ beyond which belief ascriptions may not transgress: BA only if the subject matter of the proposition expressed by A is part of $\mathbf{d}$.

Let us now see how entailments between belief ascriptions, say from BA to BC, relate both to the subject matter and the syntactic construction of A and C.

Lemma 3.2. BA *entails* BC *only if* $sm^+(C) \sqsubseteq sm^+(A)$ *in every simple doxastic model* $\mathcal{M}$.

Proof: Consider a model $\mathcal{M} = \langle S, \sqsubseteq, D, V^+, V^- \rangle$ for which $sm^+(C) \not\sqsubseteq sm^+(A)$. Given definition 6, A and C must be B-free sentences and so $|A|^{+*}$ and $|C|^{+*}$ do not depend on the choice of D. Now let $\mathcal{M}_A$ be just like $\mathcal{M}$ but with $|A|^{+*}$ in place of D. Then $sm^+(C) \not\sqsubseteq sm^+(A)$ in $\mathcal{M}_A$ and, given closure, $sm^+(C) \in |C|^{+*}$ but $sm^+(C) \notin |A|^{+*}$ (else $sm^+(C) \sqsubseteq sm^+(A)$) and so $|C|^{+*} \not\leq |A|^{+*}$. But then $\mathcal{M}_A \nvDash \mathrm{B}C$ and so $\mathrm{B}A$ does not entail $\mathrm{B}C$. $\blacksquare$

Lemma 3.3. *For* B*-free sentences* A, C *and simple doxastic models* $\mathcal{M}$:

(i) $sm^{\pm}(A) \sqsubseteq sm^{\pm}(C)$ *in each* $\mathcal{M}$ *iff* $lett(A) \subseteq lett(C)$

(ii) $sm^+(A) \sqsubseteq sm^+(C)$ *in each* $\mathcal{M}$ *iff* $lett^+(A) \subseteq lett^+(C)$ *&* $lett^-(A) \subseteq lett^-(C)$

Proof: For the left-to-right directions, we use a *canonical model* construction. Let L be the set of all literals and $D \subseteq 2^L$ (i.e. a set of sets of literals). The D-*canonical doxastic model* is $\mathcal{M}^D = \langle 2^L, \sqsubseteq, D, V^+, V^- \rangle$, where $\sqsubseteq$ is $\subseteq$ restricted to 2^L, $V^+ = \{\{p\}\}$, and $V^- = \{\{\neg p\}\}$. Given lemma 2.1, in $\mathcal{M}^D$:

$$sm^{\pm}(A) = \bigcup_{p \in lett(A)} \{p, \neg p\} \qquad sm^+(A) = \bigcup_{p \in lett^+(A)} \{p\} \cup \bigcup_{p \in lett^-(A)} \{\neg p\}$$

Now for (i), assume that $sm^{\pm}(A) \sqsubseteq sm^{\pm}(C)$ in every simple doxastic model and that $p \in lett(A)$. Then $sm^{\pm}(A) \subseteq sm^{\pm}(C)$ in $\mathcal{M}^D$, hence $\{q, \neg q \mid q \in lett(A)\} \subseteq \{q, \neg q \mid q \in lett(C)\}$, and so $p \in lett(C)$. Similar reasoning for (ii) establishes that $lett^+(A) \subseteq lett^+(C)$ and $lett^-(A) \subseteq lett^-(C)$.

For (i) right-to-left: if $lett(A) \subseteq lett(C)$ then, for an arbitrary simple doxastic model $\mathcal{M}$:

$$\bigsqcup_{p \in lett(A)} sm^+(p) \quad \sqsubseteq \quad \bigsqcup_{p \in lett(C)} sm^+(p)$$

in $\mathcal{M}$ and so, by lemma 2.1, $sm^+(A) \sqsubseteq sm^+(C)$ in $\mathcal{M}$, and hence (since $\mathcal{M}$ was arbitrary) in every simple doxastic model. The reasoning for (ii) right-to-left is similar. $\blacksquare$

Theorem 3.4. BA *entails* BC *only if* $lett(C) \subseteq lett(A)$.

Proof: Assume BA entails BC and $p \in lett(C)$. Then $sm^+(C) \sqsubseteq sm^+(A)$ in every simple doxastic model (lemma 3.2), hence $lett^+(A) \subseteq lett^+(C)$ and $lett^-(A) \subseteq lett^-(C)$ (lemma 3.3). Moreover, either $p \in lett^+(C)$, in which case $p \in lett^+(A)$, or else $p \in lett^-(C)$, in which case $p \in lett^-(A)$. Either way, $p \in lett(A)$, and so $lett(C) \subseteq lett(A)$. $\blacksquare$

This is the relevant logician's *variable sharing condition* in overdrive! In particular, Bp does not entail $\mathrm{B}(p \vee q)$. (This will be important in §6, when we consider the problems of logical omniscience.) In fact, the behaviour of disjunction within belief reports is fully accounted for by the equivalences listed in figure 1 (that is, BA and BC are equivalent when $A \equiv_e C$), plus the rule that, if BA entails BC, then $\mathrm{B}(A \vee B)$ entails $\mathrm{B}(C \vee B)$. Conjunction within belief reports

is characterised by a similar rule, plus conjunction elimination (from $B(A \wedge C)$ to BA). These axioms and rules give us a deductive relation on sentences, the transitive closure of which is none other than Angell's *Analytic Containment* (Angell 1989):

Definition 7 (Analytic Containment). $\vdash_{AC}$ is the smallest relation between sentences which includes the axioms and is closed under the rules shown in figure 2. We read $A \vdash_{AC} C$ as A *analytically implies* C and say that the associated proposition $|A|^{+*}$ *analytically contains* $|C|^{+*}$.

(The deductive equivalence of these rules to Angell's formulation is shown in Fine 2016.)

(EEQ) $A \vdash_{AC} C$ if $A \equiv_e C$ according to figure 1

($\wedge$E) $A \wedge B \vdash_{AC} A$

$$\frac{A \vdash_{AC} B \quad B \vdash_{AC} C}{A \vdash_{AC} C} \text{ (TRANS)}$$

$$\frac{A \vdash_{AC} B}{(A \wedge C) \vdash_{AC} (B \wedge C)} \text{ ($\wedge$I)} \qquad \frac{A \vdash_{AC} B}{(A \vee C) \vdash_{AC} (B \vee C)} \text{ ($\vee$I)}$$

Figure 2: Analytic Containment

Theorem 3.5. BA *entails* BC *iff* $A \vdash_{AC} C$.

Proof: Fine (2016, theorem 23) shows that $A \vdash_{AC} C$ iff $|C|^{+*} \leq |A|^{+*}$ in every model $\mathcal{M}$, which is the case iff BA entails BC. ∎

These results bring out an idea discussed by Yablo (2014) and Elgin (2021), that subject matter is a limiting factor on knowledge or belief ascriptions. (They both focus on the case of knowledge.) Yablo's idea (discussed in §1, now transposed to the case of belief) is that you may believe that you closed your front door, without thereby believing that any future evidence to the contrary is misleading. For your belief concerns the subject matter *how the door is*, which need not include the subject matter *whether there is evidence from other sources concerning how the door is* (Yablo 2014, §7.4).

Elgin (2021) argues, on independent grounds, that knowledge is closed under known analytic implication. That is, knowing both that A and that A analytically implies C implies knowing that C. The connection to subject matter is given by the truthmaker semantics for analytic implication, on which analytic implication ($A \vdash_{AC} C$) amounts to conjunctive parthood ($|C|^{+*} \leq |A|^{+*}$). Note that, on the

present semantics, we have no way to express analytic implication, and hence no way to express belief or knowledge of an analytic implication, in the object language.

4   Embedded Belief

The account presented in §3 does not allow us to embed belief reports. We cannot express, and so cannot analyse, an agent who believes she believes some proposition. The restriction is due to our semantics. Belief reports are analysed in terms of propositions (sets of states of affairs). So to analyse an agent who believes she believes some p, BBp, we first need to say what proposition Bp expresses. This is turn requires saying which states s make Bp true. At present, we can say whether or not Bp is true on a model, but not what *makes* it true or false. We have also focused on the beliefs of a single agent only, whereas much of the interest in epistemic logic arises in the multi-agent case. Let us now remedy these shortcomings.

We consider a finite number of agents i, adding a belief operator B$_i$ to the propositional language for each, so that B$_i A$ is a sentence whenever A is. Semantically, we introduce for each agent i a partial function δ_i from states s to pairs of regular sets of states $\langle D^+, D^- \rangle$, understood as the bilateral proposition giving agent i's doxastic state according to s. D^+ gives the agent's *positive* doxastic state and D^- the agent's *negative* doxastic state, relativised to a state of affairs s.

Definition 8. A *model for n agents* is a quintuple $\mathcal{M} = \langle S, \sqsubseteq, \{\delta_i\}_{i \le n}, V^+, V^- \rangle$, where $\langle S, \sqsubseteq, V^+, V^- \rangle$ is as before and each δ_i is a partial function $S \rightharpoonup (2^S \times 2^S)$ from states to pairs of regular sets of states. Where $\delta_i s = \langle D^+, D^- \rangle$, we write $\delta_i^+ s$ for D^+ and $\delta_i^- s$ for D^-.

For agent i to believe A in virtue of s is then for A's positive content to be part of $\delta_i^+ s$:

$$s \Vdash \mathsf{B}_i A \text{ iff } |A|^{+*} \le \delta_i^+ s$$

or, spelling out the clause in full:

$$s \Vdash \mathsf{B}_i A \text{ iff (B-Up) } u \in |A|^{+*} \text{ only if } t \in \delta_i^+ s \text{ for some } t \sqsupseteq u; \text{ and}$$
$$\text{(B-Down) } t \in \delta_i^+ s \text{ only if } u \in |A|^{+*} \text{ for some } u \sqsubseteq t$$

We may simplify (B-Up):

Lemma 4.1. *(B-Up) is equivalent to (B-Up′):* $u \Vdash A$ *only if* $u \sqsubseteq \bigsqcup \delta_i^+ s$.

Proof: Given lemma 3.1(v), (B-Up) simplifies to: $u \in |A|^{+*}$ only if $u \sqsubseteq \bigsqcup \delta_i^+ s$. If $u \Vdash$ then $u \in |A|^{+*}$ and so, given (B-Up), $u \sqsubseteq \bigsqcup \delta_i^+ s$. Thus (B-Up) implies (B-Up′). Now assume (B-Up′). Then $u \in |A|^+$ implies $u \sqsubseteq \bigsqcup \delta_i^+ s$ and so $\bigsqcup |A|^+ \sqsubseteq \bigsqcup \delta_i^+ s$. But $\bigsqcup |A|^+ = \bigsqcup |A|^{+*} = sm^+(A)$ and so $u \sqsubseteq sm^+(A) \sqsubseteq \bigsqcup \delta_i^+ s$ for any $u \in |A|^{+*}$. Thus (B-Up′) implies (B-Up). ∎

We also need to say when a state makes a belief ascription false. This is where the bilateral semantics earns its keep. The negative component of a total belief state, $\delta_i^- s$, is intended to model what agent i fails to believe *in virtue of state s*. There may in general be many reasons for a failure of belief, from explicit evidence to the contrary to the agent's lack of interest in the topic. We reflect this by allowing a state to be a falsitymaker for $B_i A$ independently of whether it is a truthmaker for $B_i A$. For agent i to fail to believe A in virtue of s, is for A's positive content to be included in $\delta_i^- s$:

$$s \dashv B_i A \text{ iff } |A|^{+*} \subseteq \delta_i^- s$$

As before, we understand exact entailment ($A \vDash_e B$) and exact equivalence ($A \equiv_e B$) in terms of the corresponding regular contents: $|B|^{+*} \subseteq |A|^{+*}$ and $|A|^{+*} = |B|^{+*}$, respectively. We shall henceforth call this the *truthmaker semantics (TMS) account of belief ascription*.

The falsitymaking clause uses (set-theoretic) inclusion, $|A|^{+*} \subseteq \delta_s^-$, where the truthmaking clause uses parthood, $|A|^{+*} \leq \delta_i^+ s$. Why is this? Suppose that, in virtue of some state s, I do not believe the object to be red. We might take s to be a state of visual evidence pertaining to the object's colour. Then s must thereby count as evidence that the object is not scarlet. In forming my beliefs rationally based on that evidence, s makes it the case that I do not believe the object to be scarlet. The relationship between those contents, *it is scarlet* and *it is red*, is one of *inclusion*, not of parthood: any truthmaker for 'it is scarlet' is thereby a truthmaker for 'it is red', but some truthmakers for 'it is red' (such as the state *that it is maroon*) will not contain any truthmaker for 'it is scarlet'.

More generally, we understand s to be a falsitymaker for BA in terms of $|A|^{+*}$'s inclusion in δ_s^-. Thus δ_s^- may be seen as the disjunction of all not-believed content, just as δ_s^+ may be seen as the conjunction of all believed content. As a consequence, $\neg B_i A$ will be an exact consequence of $\neg B_i C$ whenever $|A|^{+*} \subseteq |C|^{+*}$. In particular, $\neg B_i(A \vee C)$ will exactly entail $\neg B_i A$. (Indeed, if we think of the proposition *it is red* as the disjunction *it is either scarlet or maroon or ...*, then this entailment explains the previous example.)

By contrast, $\neg B_i A$ will not exactly entail $\neg B_i(A \wedge C)$. Because the object appears blue (s), I do not believe it to be red; and because the object appears cylindrical (u), I do not believe it to be rectangular. s alone explains my lack of belief that it is red, whereas s and u together explain my lack of belief that it is both red and rectangular. Of course, given that the object appears blue, I may infer and so come to disbelieve that it is not both red and rectangular. But then I lack that belief in virtue of s combined with my inference (t), not in virtue of s alone.

For positive belief reports $B_i A$, we retain the *if* direction of theorem 3.5:

Theorem 4.2. $B_i A \vDash_e BC$ *if* $A \vdash_{AC} C$.

Proof: Suppose $A \vdash_{AC} C$. From Fine 2016, theorem 23, $|C|^{+*} \leq |A|^{+*}$, so for any t: $|C|^{+*} \leq \delta^+ t$ if $|A|^{+*} \leq \delta^+ t$. Then we have $|B_i C|^+ \subseteq |B_i A|^+$, hence $|B_i C|^{+*} \subseteq |B_i A|^{+*}$ and so $B_i A \vDash_e BC$. ∎

In particular, $B_i(A \wedge C)$ exactly entails $B_i A$ but $B_i A$ does not exactly entail $B_i(A \vee C)$. Notice that we have:

$$B_i(A \wedge C) \vDash_e B_i A \quad \text{but} \quad \neg B_i A \nvDash_e \neg B_i(A \wedge C)$$
$$\neg B_i A \vDash_e \neg B_i(A \vee C) \quad \text{but} \quad B_i A \nvDash_e B_i(A \vee C)$$

Exact entailment does not contrapose. This not a specific feature of doxastic models, for even on the class of basic models $\langle S, \sqsubseteq, V^+, V^- \rangle$, $A \vDash_e C$ does not imply $\neg C \vDash_e \neg A$.

We will look at more belief ascription entailments in §6, when we discuss logical omniscience.

5 Truth and Introspection

Modal epistemic logics allow us to relate familiar axioms governing the belief operator to principles of the accessibility relation R on worlds. Of particular interest are the T and B axioms:

(T) $B_i A \to A$ (Truth)

(4) $B_i A \to B_i B_i A$ (Positive introspection)

In the classical modal setting, these correspond to reflexivity and transitivity of the accessibility relation R, respectively. What is the picture on the TMS account?

It will be useful to introduce *exact* and *inexact* conditionals, at the level of models rather than states:

Definition 9. $A \to C$ is true in a model $\mathcal{M}$ ($\mathcal{M} \vDash A \to C$) just in case $s \Vdash C$ whenever $s \Vdash A$ in $\mathcal{M}$. Similarly, $A \twoheadrightarrow C$ is true in $\mathcal{M}$ ($\mathcal{M} \vDash A \twoheadrightarrow C$) just in case $s \Vdash C$ whenever $s \Vdash A$ in $\mathcal{M}$. $A \to C$ or $A \twoheadrightarrow C$ is *valid* on a class of models $\mathcal{C}$ when it is true in all models in $\mathcal{C}$.

The exact truth axiom, $B_i A \twoheadrightarrow A$, is of little interest, for what makes it true that an agent has a given belief will not in general be what makes that belief true. The inexact version, (T) above, adequately captures the restriction to true belief. It says that states which contain a truthmaker for $B_i A$ also contain a truthmaker for A. If we understand *truth-at-a-state* in terms of that state containing an appropriate truthmaker, then it implies that, at any state, the beliefs held there are true. The semantic condition required for (T) is just the projection function analogue of reflexivity:

(INCLUSION) $s \in \delta_i^+ s$

Lemma 5.1. $B_i A \to A$ *is valid on the class of models which satisfy* INCLUSION *(for each agent i and state s).*

Proof: Assume $\mathcal{M}$ satisfies (INCLUSION) and $s \Vvdash B_i A$. Then for some $s^- \sqsubseteq s$, $s^- \Vdash B_i A$ and so $|A|^{+*} \le \delta_i s^-$. By (INCLUSION), $s^- \in \delta_i^+ s^-$ and so, by (B-Down), $u \in |A|^{+*}$ for some $u \sqsubseteq s^-$. Since $u \in |A|^{+*}$, there is some $u^- \sqsubseteq u$ for which $u \Vdash A$. Then $u^- \sqsubseteq u \sqsubseteq s^- \sqsubseteq s$ and so $s \Vvdash A$. ∎

Now we turn to positive introspection. The picture here is much more complex than in the classical setting, which considers just truth at accessible worlds. In the truthmaker setting, by contrast, we must consider the propositions P contained by a set of accessible states, $P \le \delta_i s$. The condition required to guarantee (4) is the following, for any content P and state u:

($\star$) If $P \le \delta_i^+ s$ then (i) $s \sqsubseteq \mathbf{p}$ and (ii) if $t \in \delta_i^+ s$ then $P \le \delta_i^+ u$ for some $u \sqsubseteq t$.

Lemma 5.2. $B_i A \twoheadrightarrow B_i B_i A$ and $B_i A \to B_i B_i A$ *are both valid on the class of models which satisfy* ($\star$).

Proof: We first show that $B_i A \twoheadrightarrow B_i B_i A$ is true on all such models. Assume $\mathcal{M}$ satisfies ($\star$) and $s \Vdash B_i A$. Then $|A|^{+*} \le \delta_i^+ s$. Let $\mathbf{a} = \bigsqcup |A|^{+*}$ and $\mathbf{ba} = \bigsqcup |B_i A|^{+*}$. We show $|B_i A|^{+*} \le \delta_i^+ s$:

(Up): Consider any $t \Vdash B_i A$. Then $|A|^{+*} \le \bigsqcup \delta_i^+ t$ and, given ($\star$i), $t \sqsubseteq \mathbf{a}$. Since this holds for all $t \in |B_i A|^+$, we have $\bigsqcup |B_i A|^+ = \bigsqcup |B_i A|^{+*} = \mathbf{ba} \sqsubseteq \mathbf{a}$. Given $|A|^{+*} \le \delta_i^+ s$, we also have $\mathbf{a} \sqsubseteq \bigsqcup \delta_i^+ s$. Now consider any $u \in |B_i A|^{+*}$. Then $u \sqsubseteq \mathbf{ba}$, hence $u \sqsubseteq \mathbf{ba} \sqsubseteq \mathbf{a} \sqsubseteq \bigsqcup \delta_i s$ and so $u \sqsubseteq \bigsqcup \delta_i^+ s$.

(Down): Assume $t \in \delta_i^+ s$. From ($\star$ii), $|A|^{+*} \le \delta_i^+ u$ for some $u \sqsubseteq t$. Then $u \Vdash B_i A$ and so $u \in |B_i A|^{+*}$.

It follows that $|B_i A|^{+*} \le \delta_i^+ s$ and so $s \Vdash B_i B_i A$. Thus $B_i A \twoheadrightarrow B_i B_i A$ is valid in models satisfying ($\star$). For the inexact case, assume $s \Vvdash B_i A$. Then $s^- \Vdash B_i A$ for some $s^- \sqsubseteq s$. By the previous reasoning, $s^- \Vdash B_i B_i A$ and so $s \Vvdash B_i B_i A$. So $B_i A \to B_i B_i A$ too is valid in models satisfying ($\star$). ∎

Condition ($\star$) does not look particularly intuitive, especially when compared to its classical analogue, transitivity of R. What does ($\star$) mean? Part (i) is a restriction on subject matter: if state s is to make it true that the agent believes that A, then s must be wholly relevant to A's subject matter. (That is not to say that it makes A true: it need not.) Part (ii) says that for any containment relationship $P \le Q$, each member t of Q has a part u for which $P \le \delta_i^+ u$. (This is a slight generalisation, for ($\star$) applies only when $Q = \delta_i^+ s$ for some s. But the point here is that (ii), unlike (i), is not primarily about state s.) There is a *kind* of transitivity here: if we can go from contents P to Q, and from Q to some $\delta_i^+ u$ (with u as above), then we can go directly from P to $\delta_i^+ u$.

6 Logical Omniscience

A logic of belief tells us that, as a matter of logic, if certain things are believed then so must be some other things. If the relationship between those things is

some strong notion of logical consequence, then we may have a problem: agents will be said to believe far more than any real agent could. Agents will be treated as believing all consequences of what they believe, including all logical validities. Someone with inconsistent beliefs – all of us – might even be treated as believing everything. This is the *problem of logical omniscience*. What does the TMS account have to say for itself in this regard and how does it compare to other responses in the literature?

On closer inspection, there is not just one problem of logical omniscience here. Fagin et al. (1995, 335–6) and van Ditmarsch et al. (2008, 23) discuss the following closure conditions, all of which have been questioned:

(C1) If $B_i A$ and A entails C, then $B_i C$

(C2) If A is valid, then $B_i A$

(C3) $\neg(B_i A \wedge B_i \neg A)$ is valid

(C4) $B_i A$ and $B_i C$ are equivalent whenever A and C are

(C5) If $B_i A$ and $A \rightarrow C$ is valid, then $B_i C$

(C6) If $B_i A$ and $B_i (A \rightarrow C)$, then $B_i C$

(C7) $B_i (A \wedge C)$ iff $B_i A$ and $B_i C$

(C8) If $B_i A$, then $B_i (A \vee C)$

Problems of logical omniscience are often addressed by weakening the logic underlying worlds (as, e.g., in Hintikka 1975, Levesque 1984, and Wansing 1990). If a paraconsistent logic is adopted, for example, then it is easy to model agents with inconsistent beliefs. And for any choice of sub-classical logic, it will be easy to model agents who do not believe all classical consequences of what they believe.

It is often undesirable to weaken the operative notion of logical consequence in this way, however (Fagin and Halpern 1988; Jago 2007). For although we want to accommodate an agent with inconsistent *beliefs*, we likely want our logic of belief *ascription* to remain consistent and to preserve some classical (or other strong) notion of consequence. One way to achieve this (within a worlds framework) is to distinguish between two classes of worlds. The *normal* (or *possible*) worlds behave classically, whereas the *non-normal* (or *impossible*) worlds need not. Belief is defined over all worlds, so that an agent's beliefs need not be closed under classical consequence. But consequence for the logic is defined over the normal worlds only, so that classical consequence is preserved. This approach is not without problems (Jago 2014a; Berto and Jago 2019), but that is not our central concern here.

Let us see how the TMS account compares. As already noted, identity for TMS propositions obeys the commutativity, associativity, distributivity, idempotence, De Morgan, and double negation laws. We therefore have the corresponding exact equivalences for belief ascriptions: $B_i A \equiv_e B_i C$ whenever $|A|^{+*} = |C|^{+*}$ (i.e. those equivalences $A \equiv_e C$ given in figure 1). We also have closure under

conjunctive parthood: $B_i A$ exactly entails $B_i C$ whenever $|C|^+ \leq |A|^+$. In particular, $B_i(A \wedge C)$ exactly entails $B_i A$. And conversely, beliefs are closed under conjunction introduction:

Lemma 6.1. $B_i A, B_i C$ *together exactly entail* $B_i(A \wedge C)$.

Proof: Immediate from lemma 3.1(iv) and the clause for B_i. ∎

A key feature of the TMS approach is that belief is not closed under disjunction introduction. $B_i A$ does not exactly entail $B_i(A \vee C)$ for arbitrary C, for $|A \vee C|^+$ is not in general a conjunctive part of $|A|^+$. The TMS approach validates only (C7) in our list of closure principles, where equivalence is understood classically. (It also validates the analogue of (C4) with equivalence understood as exact equivalence.)

Note that disjunction introduction *is* a valid exact entailment: A exactly entails $A \vee C$ for arbitrary C. So belief is not closed under exact entailment (and so not under any of the stronger notions of entailment) on the TMS approach. Failures of omniscience are not achieved by weakening the operative notion of entailment, in other words. Indeed, TMS can recapture a wide variety of consequence relations, including classical entailment. We might take our operative notion of entailment to be the classical one, whilst still avoiding (C8), the entailment from $B_i A$ to $B_i(A \vee C)$. And, unlike on the non-normal worlds approach described above, TMS does this with a *uniform* treatment of the connectives, thus avoiding the compositionality objection.

The TMS approach thus takes beliefs to be closed under conjunction but not under disjunction. It is, in this respect, a dual approach to Lewis's (1982) *logic for equivocators*, on which beliefs are closed under disjunction but not conjunction: $B_i A$ and $B_i C$ together do not imply $B_i(A \wedge C)$. Lewis considers failures of omniscience which are due to a *fragmented* system of beliefs:

> My system of beliefs was broken into (overlapping) fragments. Different fragments came into action in different situations, and the whole system of beliefs never manifested itself all at once. (Lewis 1982, 436)

Fagin and Halpern (1988) give a formal model of belief along these lines, which they call a model of 'local reasoning'. They think in terms of multiple 'frames of mind', each corresponding to one of Lewis's belief fragments. On either approach, an agent believes whatever is believed in *some* fragment. This results in a degree of inconsistency tolerance. An agent may believe that A in one fragment and that $\neg A$ in another, and so believe both overall, without thereby believing arbitrary propositions. They need not believe the explicit contradiction, $A \wedge \neg A$, if they never 'put two and two together' and combine these fragments of belief.

On Lewis's and Fagin and Halpern's approach, each fragment of belief is treated in the classical way, as a set of possible worlds. Thus each fragment must be internally consistent, must contain every classical tautology, and must be closed under logical consequence. So even though the agent's beliefs as a whole are not classically closed, it remains the case that they are closed under single-premise entailment. Thus, Lewis's approach satisfies (C1), (C2), (C4), (C5), and

(C8), plus (C7) left-to-right. As a response to the logical omniscience problem, Lewis's approach leaves a lot to be desired. It nevertheless captures a degree of psychological realism. Agents often fail to combine individual beliefs they hold. Believing that A and separately that C does not automatically produce the combined belief that $A \wedge C$.

We can easily incorporate Lewis's insight into the TMS approach. We view a doxastic agent as a fragmented system, as with Lewis, but with each local fragment modelled as on the TMS account. Thus, a *TMS model of local reasoning* for a single agent with multiple frames of mind $1, \ldots, n$ is a multi-agent model as given in definition 8, with each doxastic function δ_i now understood as giving the content of fragment i of the agent's total belief state. We read $B_i A$ as 'the agent believes that A in fragment i' and we introduce a new modality, B^+, to capture her beliefs overall, just as Lewis does:

$$s \Vdash B^+ A \text{ iff for some } i : s \Vdash B_i A$$

To model *multiple* fragmented agents like this, we simply partition the fragments $1, \ldots, n$ into equivalence classes $E_1, \ldots, E_m$, each with its own belief modality B_k^+, with the slightly amended clause:

$$s \Vdash B_k^+ A \text{ iff for some } i \in E_k : s \Vdash B_i A$$

In what follows, I will ignore this complication with multiple fragments of belief and focus on agents with a single, unfragmented belief state.

Lewis's approach is both a model of and a reasonable explanation for (C7)'s failure from right-to-left. The TMS account is a model of (C8)'s failure. Can it also supply a reasonable explanation? This will be the topic of the next section.

7 Concept Possession

A plausible explanation of why (C8) fails is that it seems inadmissible to ascribe to an agent a belief that goes beyond her conceptual repertoire. Stalnaker (1984, 88) gives the example of William III, who in 1700 believed that war with France could be avoided. Yet it seems we cannot say that William thereby believed that *nuclear* war could be avoided, even though this is clearly implied by what he believed, for he had no grasp of the concept of nuclear war. Similarly, we should not say that he believed that either war or nuclear war could be avoided, even though that proposition is equivalent to what he did believe. Since a disjunction $A \vee B$ may well involve concepts not present in A alone, we appear to have a plausible explanation for the failure of (C8).

Whether this explanation is acceptable depends on what it is to *possess* a concept (or, perhaps, on what concepts themselves are). On one approach, concepts are constituents of mental representations or thoughts. On the competing approach, concepts are better understood as *abilities*. In this section, I shall argue as follows. The former understanding does not provide a good explanation of (C8)'s failure.

The latter might, but has a serious problem, which itself may be resolved by reformulating it in terms of truthmaker semantics.

The former cluster of approaches understand concepts as something like items of mental vocabulary or components of inner representations. This is an idea we find in defenders of the *representational theory of mind*, such as Carruthers (2006), Fodor (2003), and Millikan (2000). We may (for current purposes) also include here views on which concepts are abstract *Fregean senses* (Peacocke 1992). What these approaches have in common is that one may possess the concept HESPERUS without possessing the concept PHOSPHORUS, or possess EYE DOCTOR without possessing OPHTHALMOLOGIST. One can then explain the *cognitive significance* of 'Hesperus is Phosphorus' or 'ophthalmologists are eye doctors' in those terms. On such views, it is natural to take belief ascriptions to be accurate when they correspond to the agent's inner representations or Fregean thoughts. So possession of the relevant concepts (so understood) will be a necessary condition on belief ascription.

In the epistemic logic literature, *logics of awareness* draw on ideas along these lines. The approach, beginning with Fagin and Halpern 1988, is to combine a possible worlds approach with a *syntactic filter*, which specifies the primitive concepts of which the agent is aware. On Fagin and Halpern's (1988) basic logic of awareness, valuations are restricted to the sentence letters in the agent's *awareness set*. If a sentence A contains a letter p not in the awareness set, A will receive no truth-value at any epistemically accessible world. In particular, there may be accessible worlds relative to which A but not $A \vee B$ is true and hence (C8) is invalidated. The syntactic awareness filter, central to these technical accounts, may be understood as corresponding to the concepts she possesses, in the representational or Fregean sense just given.

Even setting aside technical issues, this approach has philosophical problems. Consider Anna, describing in detail her poor eye health to Cath, an ophthalmologist. Anna does not possess (in the above sense) the concept OPHTHALMOLOGIST, instead thinking of Cath as the doctor at the hospital with expertise in eye problems. We might describe the situation as follows:

(1) Anna believes that Cath is an ophthalmologist, although she wouldn't put it like that.

This ascription explains why Anna is talking to Cath, whilst conveying also that Anna does not conceptualise her (in the above sense) under OPHTHALMOLOGIST. The ascription is certainly coherent. So it is coherent to ascribe beliefs involving concepts which the agent does not herself use to represent the situation in question. But if so, concept possession (so understood) is not necessary for belief ascription and so we lose our explanation of why (C8) fails.

On the alternative understanding, concepts are understood as *abilities* of a certain kind (Dummett 1993). To possess a concept is to able to draw certain distinctions in the world. For Yalcin (2018), 'a concept determines a matrix of distinctions' such that

> To possess a concept is to have an ability to cut logical space in a certain way,

to distinguish possibilities in terms of the sorts of things that answer to the concept. (Yalcin 2018, 36)

In the example, Anna can distinguish ophthalmologists from non-ophthalmologists. What she lacks is competence with the word 'ophthalmologist' (and perhaps lacks corresponding inner representations). She stands in contrast in this respect to someone unable to distinguish doctors from non-doctors, to whom we should not attribute a belief like (1), even with the caveat. On this approach, the proper explanation of (C8)'s failure is that an agent may be competent with the concepts involved in A but not with those in B and hence not with all of those in $A \vee B$.

There is an issue here, however. The ability to distinguish Fs is just the ability to distinguish $F \wedge (F \vee G)$s, since necessarily, these are the same individuals. But clearly, the ability to distinguish Fs does not imply the ability to distinguish $F \vee G$s. Someone who cannot distinguish Gs at all will be bad at distinguishing $F \vee G$s when the Gs are prevalent. But then we must allow in general that conceptual ability with a concept $F \wedge G$ does not imply conceptual ability with G. So if competence with a concept is necessary for belief and hence for knowledge, this will imply that it is possible to know that something is an $F \wedge G$ without knowing it to be a G. We surely want to avoid results like this.

My suggestion is that we understand *conceptual competence*, of the kind required to ascribe to an agent a belief with that conceptual content, in terms of her ability to identify states of the world which correspond *exactly* to that content. Competence with a concept F amounts to the ability to discern, for any suitable x, a state which exactly decides whether x is F. That state should be an exact truthmaker for $Fx \vee \neg Fx$. The requirement of *exactness* allows for the distinction between F and $F \wedge (F \vee G)$, which the Lewis-Yablo-Yalcin possible worlds approach does not, and hence avoids the problem just discussed.

(It is an interesting question whether there is independent reason for thinking one can be competent with a conjunctive concept but not with its conjuncts. A referee suggests the following case: a creature unable to detect objects smaller than elephants might thereby count as competent with PINK ELEPHANT but not with PINK. Perhaps that is right for *some* notion of conceptual competence but, as just noted, this cannot be the operative notion if we wish to take conceptual competence as necessary for belief. One available response is as follows. The operative notion of conceptual competence should be understood relative to the agent's relevant non-conceptual abilities. Thus, an agent's competence with a visual concept like PINK amounts to the ability to discern, amongst the kind of thing she can see clearly, states which exactly decide whether each such thing is pink. There is clearly much more to be said here, however.)

This approach is clearly close to Yalcin's. There is a further similarity. For Yalcin, 'a concept determines a matrix of distinctions' (2018, 36), understood as a partition on possible worlds. Within the Lewis-Yablo tradition, partitions on worlds are also understood as explicating *subject matters*. So concepts, for Yalcin, are a certain kind of subject matter. On the TMS understanding, by contrast, subject matters are understood as 'flattened' propositions ($\mathbf{p} = \bigsqcup P$ for a unilateral proposition P: see §2). Each subject matter is a state, not a partition on states (or

worlds). We may also align concepts with subject matters, so understood. The concept F is the fusion, for every x, of all states which decide whether x is F. This is the subject matter of the proposition which says, of each individual, that it either is or isn't an F. If we further identify universally quantified propositions with infinite conjunctions over all individuals (as van Fraassen (1969) does), then the concept F will be identified with the subject matter of $\forall x(Fx \lor \neg Fx)$. On this view, as on Yalcin's account, concepts are not literally components of the propositions or beliefs which involve those concepts (except in the special case of the proposition that everything either has or lacks the concepts in question).

There is clearly much more to be said about Yalcin's approach and its implementation within the TMS account. I should note that the basic TMS account is not *committed* to this understanding of concepts. Nevertheless, the overall approach seems to me a very promising way both of explaining the failure of (C8) within the TMS account of belief and of giving an account of concept possession which avoids the worry for Yalcin discussed above.

8 The Rationality Objection

I want now to address a worry for the TMS approach to belief. It is based on the following principle:

(LEM) Rational agents believe each instance of $A \lor \neg A$.

Any rational agent should believe all instances of $A \lor \neg A$ (or nearly all instances: we may ignore those cases in which LEM is thought to be questionable). All one need do, to believe any given instance $A \lor \neg A$, is to identify that the sentence has the syntactic form it does. I needn't know what a babirusa is in the slightest to know that either babirusas are mammals or they aren't. (And if I know it, I believe it.) A similar point can be made with a no-contradiction principle:

(NC) Rational agents believe each instance of $\neg(A \land \neg A)$.

I needn't know what a babirusa is to know that they aren't both mammals and not mammals. On the TMS account, (NC) has the same effect on belief states as (LEM) because $\neg(A \land \neg A)$ is exactly equivalent to $A \lor \neg A$.

The objection is that accepting (LEM) (or (NC)) will render any restriction based on concept possession, of the kinds discussed in §7, null and void. For in believing each instance of $A \lor \neg A$, the agent is thereby treated as possessing, or being competent with, each concept. If the only reason we cannot infer from $B_i A$ to $B_i(A \lor C)$ is because of some such restriction, then we may be forced to say the agent who believes both A and $C \lor \neg C$ must thereby believe that $A \lor C$.

This objection affects the logic of awareness (§7) deeply. On Fagin and Halpern's (1988) account, being aware of the concepts in a primitive sentence p amounts to believing $p \lor \neg p$. For complex sentences A, awareness of the concepts in A amounts to awareness of p for each subsentence p of A. So an agent aware of each A will believe any valid B and, more generally, will be fully logically

omniscient in the logic of awareness (Fagin and Halpern 1988, 47, proposition 3.1).

Whether the worry also affects the TMS account depends on how we interpret (LEM). Suppose we take it to mean the following:

(LEM-e) For every state s: $s \Vdash B_i(A \vee \neg A)$

This corresponds to the following semantic condition, which leads to troubling results.

(LEM′) $(|A|^{+*} \cup |A|^{-*})^* \leq \delta_i^+ s$ for every sentence A, agent i, and state s on which δ is defined.

Lemma 8.1. *Let* $\blacktriangle = \bigsqcup\{s \mid s \Vdash A \text{ for some } A\}$. *For any model M which satisfies* (LEM′):

(i) $\bigsqcup \delta_i^+ s \sqsupseteq \blacktriangle$ *for any* s

(ii) $|A|^{+*} \leq \delta_i^+ s$ *iff for any* $t \in \delta_i^+ s$ *there is some* $u \sqsubseteq t$ *such that* $u \in |A|^{+*}$

(iii) $s \Vdash B_i A$ *iff* $u \Vdash A$ *for all* $u \in \delta_i^+ s$

Proof: (i) Consider any s for which $\delta_i^+ s$ is defined and let u by any active state. Then $u \in |A|^+$ for some A and so, by (LEM′), $u \in (|A|^{+*} \cup |A|^{-*})^* \leq \delta_i^+ s$. By (Up), $u \sqsubseteq t$ for some $t \in \delta_i^+ s$. So $\blacktriangle \sqsubseteq \bigsqcup \delta_i^+ s$.

(ii) Consider any $u \in |A|^{+*}$. By definition, $t \sqsubseteq \blacktriangle$ for any $t \Vdash A$ and so $sm^+(A) \sqsubseteq \blacktriangle$. Since $u \sqsubseteq sm^+(A)$, it follows that $u \sqsubseteq \blacktriangle$. Then by part (i), $u \sqsubseteq \bigsqcup \delta_i^+ s$ if $\delta_i^+ s$ is defined. So (Up) is trivially satisfied and hence $|A|^{+*} \leq \delta_i^+ s$ iff (Down).

(iii) Suppose $s \Vdash B_i A$ and $t \in \delta_i^+ s$. Then $|A|^{+*} \leq \delta_i^+ s$ and so there is some $u \sqsubseteq t$ such that $u \in |A|^{+*}$. Then there is a $u^- \sqsubseteq u \sqsubseteq t$ such that $u^- \Vdash A$ and so $t \Vdash A$. Now suppose $t \Vdash A$ for all $t \in \delta_i^+ s$ and consider any such t. Then there is some $u \sqsubseteq t$ such that $u \Vdash A$, hence $u \in |A|^{+*}$. Since this holds for any such t, by (ii), $|A|^{+*} \leq \delta_i^+ s$ and so $s \Vdash B_i A$.   ∎

Definition 10 (LP). $\vdash_{\text{FDE}}$ is the smallest relation between sentences containing $\vdash_{\text{AC}}$ and such that: $A \vdash_{\text{FDE}} A \vee C$. $\vdash_{\text{LP}}$ is the smallest relation between sentences (including $\top$) containing $\vdash_{\text{FDE}}$ and such that: $\top \vdash_{\text{LP}} A \vee \neg A$. We write $\vdash_{\text{LP}} A \vee \neg A$ for $\top \vdash_{\text{LP}} A \vee \neg A$.

Theorem 8.2. *Both* $B_i A \twoheadrightarrow B_i C$ *and* $B_i A \to B_i C$ *are valid on the class of models which satisfy* (LEM′) *whenever* $A \vdash_{LP} C$.

Proof: We show the result (for $\twoheadrightarrow$) holds for each axiom and rule of LP. Given theorem 4.2, it holds for each axiom and rule of AC. From lemma 8.1(iii), $s \Vdash B_i A$ iff $u \Vdash A$ for all $u \in \delta_i^+ s$ only if $u \Vdash A \vee C$ for all $u \in \delta_i^+ s$ iff $s \Vdash B_i(A \vee C)$. So $B_i A \vDash_e B_i(A \vee C)$ and so $B_i A \twoheadrightarrow B_i(A \vee C)$. Thus the result holds for each axiom and rule of FDE. Finally, given (LEM′), $\vDash_e B_i A \vee B_i \neg A$ and so the result (for $\twoheadrightarrow$) holds for each axiom and rule of LP. Then by definition, it also holds for $\to$.   ∎

In this proof, we also established:

Corollary 8.3. *(C8) is valid on models which satisfy (LEM′).*

It is clear that the TMS account must deny (LEM-e) if it is to be plausible. Fortunately, (LEM-e) is stronger than we need to capture the thought behind (LEM), which is simply that rational agents must believe $A \vee \neg A$. (LEM-e), by contrast, says that *every state* makes it true that agents believe each instance of $A \vee \neg A$. This is not at all plausible, for what makes $\mathsf{B}_i(A \vee \neg A)$ true and what makes $\mathsf{B}_i(C \vee \neg C)$ true will often be distinct states. (LEM) requires that such states exist, but not that one does the work of all. A state will treat the agent as being rational (with respect to believing instances of $A \vee \neg A$) so long as it contains, for each instance $\mathsf{B}_i(A \vee \neg A)$, a truthmaker for that instance. Those states need not all be identical. Thus the requirement on models is just that:

(LEM-i) For every state s: $s \Vdash \mathsf{B}_i(A \vee \neg A)$

With this in place of (LEM-e), the results above are blocked and, in particular, (C8) remains invalid.

It is worth delving a little deeper into how this approach accommodates the rationality requirement in (LEM) whilst avoiding (C8). Say a state w is a *world* when $w \Vdash A \vee \neg A$ for each A and let W be the set of all worlds. Any such w will have a part u such that $u \Vdash A$ or $u \dashv\vdash A$, so that $u \in |A \vee \neg A|^{+*}$. It follows that $|A \vee \neg A|^{+*} \leq W$ for each A. But it will not be the case that W contains arbitrary disjunctions, say $|A \vee C|^{+*} \leq W$, for there are worlds (the consistent $\neg A \wedge \neg C$-worlds) with no part in $|A \vee C|^{+*}$. Now if we set $\delta_i s = W$, then $s \Vdash \mathsf{B}_i(A \vee \neg A)$ for each A. In particular, we might set $\delta_i \square = W$ so that $s \Vdash \mathsf{B}_i(A \vee \neg A)$ for each state s and each A. We thus obtain models satisfying (LEM-i) but in which $s \Vdash \mathsf{B}_i A$ does not imply $s \Vdash \mathsf{B}_i(A \vee C)$.

This demonstrates that the TMS account can treat agents as rational, in the sense of (LEM) and the equivalent (NC), whilst avoiding (C8). So the rationality objection based on (LEM), which affects awareness-style approaches, does not touch the TMS account.

9   Resources and The Problem of Rational Belief

There is one further objection to the TMS account I wish to consider. It is often said that agents fail to believe some consequences of what they believe simply because they lack the resources to derive those consequences (Konolige 1986; Jago 2009). This applies to AI as well as human agents (see, e.g., Alechina et al. 2006; Jago 2006). The objection is that the TMS account fails to capture this.

Resources include time, memory, and the ability to focus on deductive reasoning. Even when we possess all the relevant concepts and combine all our beliefs pertaining to some deductive problem, complex solutions often elude us. Playing chess is a good example of this. Suppose you and I play a game of chess without time controls in which a draw counts as a win for black. This guarantees the game will have a winner. It is then a surprising mathematical fact that, at any stage of the game, one of us has a *winning strategy*: a function from the game's

previous moves to that player's next move that is mathematically guaranteed to win, regardless of how the other player plays. That strategy follows deductively from the rules and the current game position, all of which I believe. Yet *no one* has ever calculated such a strategy (except in certain endgames, where the complexity reduces significantly). In general, it's far too complicated. Even the most powerful computer doesn't play chess purely by cranking out deductive consequences of possible moves. Heuristics, rather than pure deduction, are the backbone of chess strategy.

When I make a daft move, it's not because I lack some chess-relevant concept, or because my chess-relevant beliefs fall into different fragments of my mind. It's not that the better move wasn't relevant to the question, 'how should I move next?', or somehow not part of the subject matter I'm considering. It's just that there's a limit to how much reasoning of the form, 'if I move here, you'll likely move here or here …' which I can perform before my mind gives up and starts thinking about pancakes.

Do considerations like these show that an accurate epistemic model should be built on some resource-sensitive logic? Such logics exist: *linear logic* (Girard 1987) is a prominent example. Roughly speaking, $A \vdash B$ can be derived in linear logic when A is exactly what is needed to prove B. Thus $A, A \vdash B$ may not be derivable even if $A \vdash B$ is, for the premise-list A, A represents a resource used twice. In this case, although A, A would be sufficient for B, it would be overkill in terms of resources. (There is thus a similarity with exact entailment here, in that $A \land B$ is sufficient but not exactly relevant to A's truth.)

It is not difficult to see that building an epistemic logic on top of linear logic is not the way to avoid the resourced-based problem of logical omniscience. Linear logic remains too strong in that it validates *modus ponens*, whereas agents' beliefs are not in general closed under *modus ponens*. (Just consider the rules of chess written as implications.) Linear logic is also too weak, in that the doxastic agents we are considering (let us suppose) are happy to infer from $A \vdash B$ to $A, A \vdash B$ and vice versa. Indeed, it is not obvious what it could mean for an agent to believe that A twice over, but not believe it once over. The problem is that linear logic views formulas as *resources*, whereas in epistemic logic, (embedded) formulas represent believed propositions.

This is not merely a problem for linear logic or the TMS approach. It exists *regardless* of the logic we use to model doxastic agents. The problem at heart is the tug-of-war between the deductive principles we consider constitutive of rationality and the fact that agents do not believe all consequences (given these principles) of what they believe. *Modus ponens* is constitutive of rationality if any deductive principle is. Now suppose we say, here is a chap who believes that $A \rightarrow B$ and believes that A, but he just doesn't believe that B. He's attentive to the question, he's thought about it seriously for some time, but still, he doesn't believe that B. What are we to make of this? We seem to be saying that the agent is irrational. (Perhaps there are genuine cases like this, although they are usually explained in other terms: perhaps the agent accepts some strange theory of the conditional or assigns some other meaning to '$\rightarrow$'. But in those cases, we should not say that she believes what *we* mean by $A \rightarrow B$. Anyway, typical cases are not

like this.)

This point is especially hard to avoid for those who accept the Dennettian view that the purpose of ascribing attitudes to an agent is to make rational sense, from our point of view, of her behaviour (Dennett 1987). There *may* be behaviour indicative of rejecting an instance of *modus ponens*: say, the agent's asserting both premises whilst explicitly denying the conclusion. This is evidence for the *stronger* ascription, that she believes the conclusion to be false. (And, as just noted, we might interpret such cases as her assigning an unusual meaning to the conditional.) The more typical case is where the agent seemingly fails to take any stance on the conclusion.

Elsewhere, I describe *the problem of rational knowledge* Jago (2014a;b); here I will call it *the problem of rational belief ascription*. We note that (i) rational agents seemingly believe the trivial consequences of what they believe, but (ii) they do not believe all logical consequences of what they believe. The problem is that (i) and (ii) are incompatible. Any logical consequence of a set of premises is derivable from those premises via a chain of trivial inferences and so, if one does not believe some logical consequence of what one believes, then one must fail to believe some trivial consequence of what one believes.

There is a structural similarity here with the sorites paradox, which is central to the problem of vagueness. The principle that rational agents believe the trivial consequences of what they believe plays the role that tolerance conditionals (for 'heap', say) play in the sorites. Clearly, not all such conditionals are true; but we cannot say or discover which is false. We must acknowledge the existence of borderline cases of 'heap' but we cannot say, of a particular case, that it is one such case. This is the phenomena of *unassertibility at the borderline*. Williamson (1992) puts the point by saying we can have *inexact* knowledge only and as a consequence we may not rationally *assert* precise claims in the vicinity of the borderline.

The situation is, I believe, similar in the case of belief ascriptions. We cannot rationally assert that *this* is the particular trivial inference which does not preserve the agent's belief, even if that is in fact the case. My view is that such instances of belief failure are always *indeterminate* instances of belief failure (Jago 2014a;b). There is never a case in which agent *i* determinately believes such-and-such, from which it trivially follows that *A*, such that it is determinate that *i* does not believe that *A*.

This analysis has a rather deflating consequence, however. There are *no* non-trivial logical principles of inference such that, if the determinately agent believes (or knows) the premises, then it logically follows that she determinately believes (or knows) the conclusion. That seems to imply that there can be no interesting logic of belief or knowledge. But that is too quick. A logic of belief is the attempt to use logical tools to draw interesting, non-trivial conclusions about what agents believe. Given the argument above, any such attempt *must* idealise away from the actual facts. Any interesting epistemic logic should be viewed as an idealisation away from an agent's cognitive limitations.

This does not mean we should accept the picture of *full omniscience* which accompanies the original possible worlds account. We can idealise away from

some aspects of an agent's cognitive state without thereby ignoring all of them. The TMS approach incorporates and allows us to reason about several interesting features of belief. It tolerates incompatible sets of beliefs and even internally inconsistent beliefs. It captures the idea that a total belief state is restricted to a particular subject matter. And, given the arguments of §6 and §8, it does so better than awareness-based approaches. The TMS account (unlike the semantics given in Jago 2014a;b) does not capture the way in which an agent's resources cause her beliefs to 'grey out' at some indeterminate point. But likely, neither does any genuinely useful epistemic logic.

References

Alechina, N., Jago, M. and Logan, B. (2006). Modal logics for communicating rule-based agents, *Proceedings of ECAI 06*, pp. 322–326.

Angell, R. B. (1989). Deducibility, entailment and analytic containment, in J. Norman and R. Sylvan (eds), *Directions in Relevant Logic*, Kluwer, Dordrecht, pp. 119–143.

Berto, F. and Jago, M. (2019). *Impossible Worlds*, Oxford University Press.

Carruthers, P. (2006). *The Innate Mind: Culture and Cognition*, Oxford University Press, Oxford.

Dennett, D. C. (1987). *The Intentional Stance*, MIT Press.

Dummett, M. (1993). *The seas of language*, Oxford University Press.

Elgin, S. (2021). Knowledge is closed under analytic content, *Synthese* 199: 5339–5353.

Fagin, R. and Halpern, J. (1988). Belief, awareness and limited reasoning, *Artificial Intelligence* 34: 39–76.

Fagin, R., Halpern, J., Moses, Y. and Vardi, M. (1995). *Reasoning About Knowledge*, MIT press.

Fine, K. (2016). Angellic content, *Journal of Philosophical Logic* 45(2): 199–226.

Fine, K. (2017a). A theory of truth-conditional content I: Conjunction, disjunction and negation, *Journal of Philosophical Logic* 46(6): 625–674.

Fine, K. (2017b). A theory of truth-conditional content II: Subject-matter, common content, remainder and ground, *Journal of Philosophical Logic* 46(6): 675–702.

Fine, K. (2018). Compliance and command ii, imperatives and deontics, *The Review of Symbolic Logic* 11(4): 634–664.

Fine, K. (2020). Yablo on subject-matter, *Philosophical Studies* 177: 129–171.

Fine, K. (2023). An epistemized truthmaker semantics for epistemic logic': Response to Hawke's and Özgün's 'truthmaker semantics for epistemic logic, in F. L. Faroldi and F. Van De Putte (eds), *Kit Fine on Truthmakers, Relevance, and Non-classical Logic*, Springer, pp. 337–348.

Fine, K. and Jago, M. (2019). Logic for exact entailment, *Review of Symbolic Logic* 12(3): 536–556.

Fine, K. and Jago, M. (Forthcoming). *An Introduction to Truthmaker Semantics*, Oxford University Press.

Fodor, J. (2003). *Hume Variations*, Oxford University Press, Oxford.

Girard, J.-Y. (1987). Linear logic, *Theoretical Computer Science* 50(1): 1–101.

Hawke, P. and Özgün, A. (2023). Truthmaker semantics for epistemic logic, in F. L. Faroldi and F. Van De Putte (eds), *Kit Fine on Truthmakers, Relevance, and Non-classical Logic*, Springer, pp. 295–335.

Hintikka, J. (1975). Impossible possible worlds vindicated, *Journal of Philisophical Logic* 4: 475–484.

Jago, M. (2006). *Logics for Resource-Bounded Agents*, PhD thesis, University of Nottingham.

Jago, M. (2007). Hintikka and Cresswell on logical omniscience, *Logic and Logical Philosophy* 15(4): 325–354.

Jago, M. (2009). Resources in epistemic logic, in J.-Y. Béziau and A. Costa-Leite (eds), *Dimensions of Logical Concepts*, Vol. 55, Coleção CLE, Campinas, Brazil, pp. 11–33.

Jago, M. (2014a). *The Impossible*, Oxford University Press.

Jago, M. (2014b). The problem of rational knowledge, *Erkenntnis* 79(6): 1151–1168.

Konolige, K. (1986). *A Deduction Model of Belief*, Morgan Kaufman, San Francisco.

Levesque, H. J. (1984). A logic of implicit and explicit belief, *Proceedings of the Fourth National Conference on Artificial Intelligence*, pp. 198–202.

Lewis, D. (1982). Logic for equivocators, *Noûs* 16(3): 431–441.

Millikan, R. (2000). *On Clear and Confused Ideas*, Cambridge University Press, Cambridge.

Peacocke, C. (1992). A theory of concepts.

Stalnaker, R. (1984). *Inquiry*, MIT Press, Cambridge, MA.

van Ditmarsch, H., Van Der Hoek, W. and Kooi, B. (2008). *Dynamic epistemic logic*, Vol. 337, Springer.

van Fraassen, B. (1969). Facts and tautological entailments, *Journal of Philosophy* 66(15): 477–487.

Wansing, H. (1990). A general possible worlds framework for reasoning about knowledge and belief, *Studia Logica* 49(4).

Williamson, T. (1992). Inexact knowledge, *Mind* 101(402): 217–242.

Yablo, S. (2014). *Aboutness*, Princeton University Press, Princeton, NJ.

Yablo, S. (2018). Reply to fine on "aboutness", *Philosophical Studies* 175(6): 1495–1512.

Yalcin, S. (2018). Belief as question-sensitive, *Philosophy and Phenomenological Research* 97(1): 23–47.