This is the penultimate version of an article accepted for publication by The Journal of Philosophical Logic. *Please quote only from the published version.*
The final publication will soon be available at `http://link.springer.com`

MIGHTY BELIEF REVISION

STEPHAN KRÄMER

Abstract. Belief revision theories standardly endorse a principle of intensionality to the effect that ideal doxastic agents do not discriminate between pieces of information that are equivalent within classical logic. I argue that this principle should be rejected. Its failure, on my view, does not require failures of logical omniscience on the part of the agent, but results from a view of the update as *mighty*: as encoding what the agent learns might be the case, as well as what must be. The view is motivated by consideration of a puzzle case, obtained by transposing into the context of belief revision a kind of scenario that Kit Fine has used to argue against intensionalism about counterfactuals. Employing the framework of truthmaker semantics, I go on to develop a novel account of belief revision, based on a conception of the update as mighty, which validates natural hyperintensional counterparts of the usual AGM postulates.

Belief revision theories standardly endorse a principle of *intensionality*, according to which it is a requirement of rationality on ideal doxastic agents that they do not discriminate between pieces of information that are equivalent within classical logic: whatever they are disposed to (come to or continue to) believe upon receiving the one, they are disposed to believe upon receiving the other, and vice versa. In this paper, I argue that, subject to certain qualifications, that principle should be rejected. Its argued failure does not require failures of logical omniscience on the part of the agent. It results instead from a view of the update as *mighty*—as encoding what the agent learns *might* be the case, as well as what must be.[1] Central to my argument is a puzzle case, obtained by transposing into the context of belief revision a kind of scenario that Kit Fine, in his 'Counterfactuals without Possible Worlds' (2012a, see also his 2012b), has used to argue against the principle of intensionality for (the antecedents of) counterfactuals.

[1] This sets the view defended here apart from previous approaches rejecting intensionality, which have generally been motivated by the aim of modelling less idealized doxastic agents. For various approaches of this sort, see e.g. Rantala (1982); Fagin and Halpern (1988); Jago (2014); Berto (2019); Özgün and Berto (2020).

The structure of the paper is as follows. Section 1 introduces some background assumptions and terminology and gives a more precise statement of the principle of intensionality. Section 2 describes an intensional account of rational belief revision—a form of the popular AGM approach—which is closely related to the standard possible worlds analysis of counterfactuals. Section 3 presents the puzzle cases. Section 4 applies the AGM approach to these cases and argues that it gives the wrong results. Section 5 examines and rejects some prima facie promising ways to respond to the difficulty while retaining intensionality. Section 6 introduces the basic ideas guiding my subsequent development of a *truthmaker-based*, hyperintensional approach. Section 7 introduces the conception of the update as *mighty* underlying the approach and explains how it leads to violations of the principle intensionality. Section 8 formally articulates some constraints on the rational ways of revising by mighty updates. It is shown that the account delivers the intuitively correct verdicts in the problem cases while retaining those components of the AGM account that are not undermined by those examples. Section 9, finally, describes in more general terms the advantages I take the truthmaker-based approach to offer while identifying some open questions for future research to pursue.

1. Belief Revision and Intensionality

At any given time, doxastic agents like ourselves have a set of beliefs, and they have dispositions to revise their beliefs in certain ways under certain circumstances. For brevity, we shall refer to such dispositions simply as *dispositions to revise*, and we shall refer to relevant circumstances as *occasions for revision*. The combination of a total system of beliefs and a total set of dispositions to revise we may call a *(complete) doxastic state*. Call a complete doxastic state *(ideally rationally) permissible* iff it could be the doxastic state of an ideally rational doxastic agent (short: ideal agent). We may also call a partial doxastic state permissible iff it has a permissible complete extension. The aim of a theory of belief revision, as I here conceive of it, is to capture the general, logico-structural properties that are held by any permissible complete doxastic state.

It is standard to assume that for any ideal agent and for any possible occasion for revision, the agent's dispositions to revise determine a *unique* result, i.e. a unique set of beliefs comprising all and only those beliefs the agent would hold after exercising their dispositions. The dispositions to revise of an ideal agent may then be represented by a function mapping every possible occasion for revision to a revised belief system.[2]

[2] Of course, there may be occasion to revise the new belief system again. To study the constraints on *iterated* belief revision, we should have to assume either that the initial doxastic state includes

Let us call occasions for revision *dynamically equivalent* iff no permissible doxastic state discriminates between them. That is to say, occasions for revision o_1 and o_2 are dynamically equivalent just in case for any function f representing the dispositions to revise in some permissible doxastic state, $f(o_1) = f(o_2)$. It is a standard (if often tacit) assumption that one way to characterize a sufficient condition for dynamic equivalence is in terms of a *proposition* suitably related to the occasion of revision—call this proposition the *update*. This seems quite plausible. Presumably, the rationality or otherwise of a possible response to an occasion for revision can depend only on what the doxastic agent *learns*, or what *information* they *receive*, on that occasion. If what the agent learns on occasion o_1 is the same as what they learn on occasion o_2, then rationality seems to require that the agent make the same adjustments to their beliefs in both situations. Assuming that the totality of what the agent learns can always be represented by a proposition, we may take that proposition to be the update and conclude that occasions for revision with the same update are dynamically equivalent. We shall later say more about how to make these ideas precise. For now, note that given the dynamic equivalence of situations with the same update, for the purposes of a theory of belief revision, we may identify occasions for revisions with their associated updates, and we may represent an agent's dispositions to revise as a function mapping *each possible update*[3] to a revised belief system. We shall also describe possible updates as dynamically equivalent when the associated occasions for revision are. A principle of intensionality for updates may now be stated as follows:

> **Intensionality:** For any possible updates P and Q, if P is logically equivalent to Q then P is dynamically equivalent to Q.

In this formulation, the principle presupposes a notion of logical equivalence for updates. The most common approach in the literature is to identify updates with sentences of some formal, propositional language. An alternative is to assume a notion of a logically possible world, and to identify updates with the sets of logically possible worlds in which they are true. Logical equivalence for updates is then simply the identity relation, and so adopting this kind of conception of update will automatically ensure that

dispositions to revise arbitrary belief states (or at least arbitrary ones reachable from the present belief state by some sequence of revisions) by new information, or that the dispositions to revise the initial doxastic state determine not only the new belief state but also new dispositions to revise that belief state. For the purposes of this paper, we restrict attention to singular, i.e. non-iterated belief revision.

3 By a possible update I mean a proposition which is the update in some possible situation for an ideal doxastic agent. It is a further question whether every update possible in this sense is also possibly true. For present purposes, though, we may assume that this is so.

Intensionality holds. The two approaches may be connected, relative to a chosen formal language, by identifying logically possible worlds with the corresponding maximal consistent sets of sentences of the language.[4]

2. A Possible Worlds Approach

Within a possible worlds framework, we can formulate a prima facie attractive theory of rational belief revision that is closely related to the standard possible worlds account of *counterfactuals*.[5] On this account of counterfactuals, recall, we assume that for any given possible world w, there is an *ordering* of all worlds according to their comparative similarity, in some suitable sense, to w. A counterfactual $A \mathrel{\Box\!\!\rightarrow} C$ is then taken to be true at w iff all those worlds at which A is true which are *closest*—i.e. most similar—to w are worlds at which C is also true.

Under the analogous approach to belief revision, both belief systems and updates are identified with a set of logically possible worlds. A doxastic state accordingly consists of a set B of possible worlds representing the belief system, and a function mapping any set of possible worlds P—the update—to a set of possible worlds $B * P$—the revised belief system. It is assumed that in any ideally rational doxastic state, B is non-empty. The logical constraints on the revision function are stated by appeal to an ordering on the worlds, formally similar to the similarity orderings by which counterfactuals are interpreted.[6] Informally, we may think of the ordering as representing the comparative plausibility of the worlds by the lights of the agent, or perhaps the strength with which the worlds are *excluded* or *disbelieved* by the agent. The worlds at which the agent's

4 Analogous questions of granularity may also be raised with respect to the other component of a doxastic state, i.e. the total system of beliefs. Our focus in this paper, though, will be on the intensionality or otherwise of the update.

5 The classical sources are Stalnaker (1968) and Lewis (1973).

6 The idea of basing belief revision theory semantically on an ordering of worlds is familiar in the literature. My presentation here largely follows Huber (2013). The approach based on plausibility orderings can equivalently be stated in terms of plausibility spheres, just like the Lewis/Stalnaker semantics can be stated in terms of similarity spheres instead of similarity orderings. Modulo the subtleties surrounding the condition (≤4) mentioned below, the present approach is thus equivalent to the sphere-based approach first described by Grove (1988). The same kind of ordering of worlds, under the label of *faithful assignments*, is used in Katsuno and Mendelzon (1991) to prove a representation theorem for AGM revision operations (the counterpart of (≤4) is not needed there, since the authors assume the underlying language to be based on a finite set of propositional letters). For a useful overview of equivalent characterizations of the AGM model, see chapter 4 of Fermé and Hansson (2018) and especially section 4.1, which discusses the various possible worlds based models.

beliefs are true are the most plausible ones, which are not excluded or disbelieved at all. All other worlds are excluded, but some more firmly than others, in which case they are treated as less plausible.

More formally, given a belief system B, we call a *plausibility ordering centered on B* any two-place relation $\leq$ on the worlds such that for all worlds w, v, u:

($\leq$1) $w \leq v$ or $v \leq w$
($\leq$2) if $w \leq v$ and $v \leq u$ then $w \leq u$
($\leq$3) $w \in B$ iff $w \leq z$ for all $z \in W$
($\leq$4) if $\emptyset \subset A \subseteq W$, then $\{z \in A\colon z \leq y$ whenever $y \in A\} \neq \emptyset$

Informally, $w \leq v$ means that w is at least as plausible as v. ($\leq$1)–($\leq$3) ensure that the plausibility ordering is transitive, that any two worlds are comparable in terms of their plausibility, and that all and only the members of B are maximally plausible. The final condition ($\leq$4), as we shall see, is of special importance for our purposes: it ensures that any non-empty set of worlds has a maximally plausible member. The crucial claim is now that for any ideally rational doxastic state with belief system B and revision function $*$, there exists a plausibility ordering $\leq$ of the worlds centered on B such that for every possible update P, $B * P = \{z \in P\colon z \leq y$ whenever $y \in P\}$: the revision by update P is always the set of the most plausible P-worlds.

This account of belief revision is near-equivalent to the popular AGM theory of belief revision (Alchourrón et al. (1985)).[7] Within AGM, a belief system is modelled by a set K of sentences of a propositional language $\mathcal{L}$, an update is modelled by a single sentence α from $\mathcal{L}$, and the dispositions to revise are modelled by a function mapping K and any such α to a new belief system $K * \alpha$. The theory then includes the following eight postulates to be satisfied by any ideally rational belief set and revision function (where $K + \alpha$ is the closure under logical consequence of $K \cup \{\alpha\}$):

Closure	$K * \alpha$ is closed under logical consequence
Success	$\alpha \in K * \alpha$
Inclusion	$K + \alpha \supseteq K * \alpha$
Vacuity	$K * \alpha \supseteq K + \alpha$ if $K \cup \{\alpha\}$ is consistent
Consistency	$K * \alpha$ is consistent if α is
Intensionality	$K * \alpha = K * \beta$ if α and β are logically equivalent
Superexpansion	$(K * \alpha) + \beta \supseteq K * (\alpha \wedge \beta)$
Subexpansion	$K * (\alpha \wedge \beta) \supseteq (K * \alpha) + \beta$ if $(K * \alpha) \cup \beta$ is consistent

[7] For an accessible introduction, see again Huber (2013) or Hansson (2017).

We shall sometimes refer to the last two postulates as the *supplementary* AGM postulates, and to the other six as the *basic* AGM postulates.[8]

It is known that from any AGM belief set K and revision function $*$, one can construct a possible worlds interpretation of $\mathcal{L}$ and an ordering $\leq$ on the worlds, centered on the set of worlds at which K is true, which satisfies conditions ($\leq$1)–($\leq$3) as well as a weakened version of ($\leq$4).[9] Say that a formula $\alpha \in \mathcal{L}$ *expresses* a set of worlds (under the given interpretation) iff it is true at exactly those worlds. Then the relevant weakening of ($\leq$4) says that any non-empty set of worlds *expressed by some formula in* $\mathcal{L}$ has a maximally plausible member. Conversely, given a plausibility ordering centered on B on a set of worlds W and an interpretation of $\mathcal{L}$ relative to W, one can define a corresponding AGM-style revision operator for the belief set true exactly at the members of B which satisfies the AGM postulates.[10] For most of the discussion to follow, we may treat AGM and the plausibility based possible worlds approach as equivalent, and refer to them indiscriminately as the AGM approach or the possible worlds approach.

3. Of Dominos and Matches

In this section, I shall describe some partial doxastic states and argue that they are rationally permissible, i.e. that they have complete extensions that could be the doxastic state of an ideally rational agent. In the next section I will then show that the permissibility of these doxastic states is in conflict with the AGM approach.

For definiteness, imagine a particular doxastic agent, Dom. His relevant beliefs concern an infinite sequence of domino stones, arranged like this[11]

[] [] [] [] . . .

8 In the literature they are also called the *basic* and *supplementary Gärdenfors postulates* for revision, respectively; cf. (Hansson, 2017: §3). There is some variation also in the labels for the individual postulates. *Intensionality* is sometimes called *Extensionality* and *Superexpansion* and *Subexpansion* are sometimes just called *Conjunction 1* and *2*.

9 That only the weakened version of ($\leq$4) is guaranteed is why I said the above account is *near*-equivalent to AGM. We will see below that this detail is somewhat relevant to our purposes.

10 These results are due to Adam Grove (1988).

11 The scenario is essentially identical to the first example described in Fine (2012a), except that Fine's scenario features rocks instead of domino stones. Note that our case strictly requires only that our agent *has* the relevant beliefs about domino stones, not that these beliefs are accurate. But for presentational purposes it seemed helpful to me to suppose the situation to be as the agent believes it to be.

We assume that each stone can only fall to the right, not to the left. We refer to the stones as $s_1, s_2, \ldots$, respectively, with s_1 being the leftmost stone, and s_{n+1} the stone immediately to the right of s_n. Let F_n be the proposition that stone n fell. We suppose that as a matter of fact, no stone fell.

For each n, Dom believes that $\neg F_n$. Furthermore, he has the following dispositions to revise: If Dom were to learn that F_n, then he would come to believe that F_m for all m with $m \geq n$. At the same time, he would retain the belief that $\neg F_m$ for all m with $m < n$.

It will be helpful to introduce some notation to describe the doxastic state more succinctly. Let us write $P \Rightarrow Q$ for the claim that Dom is disposed to (come to or continue to) believe that Q upon learning that P—i.e. on any occasion for revision whose update is the proposition that P.[12] Slightly artificially, we write $\Rightarrow Q$ to say that Dom believes that Q (since this is like saying that he is disposed to believe that Q upon learning nothing). We can now summarize the partial doxastic state $\mathcal{D}$ we have ascribed to Dom as follows:

(B) $\Rightarrow \neg F_n$ for all n

(D.+) $F_n \Rightarrow F_m$ for any $m \geq n$

(D.−) $F_n \Rightarrow \neg F_m$ for any $m < n$

Dom's dispositions to revise may be seen simply as reflecting an awareness of the nature of the setup as described above: since each stone can only fall to the right, knocking over every subsequent stone, if Dom learns F_n he also comes to believe F_{n+1}, F_{n+2}, ... and accordingly gives up $\neg F_{n+1}$, $\neg F_{n+2}$, ... but since each stone can only fall to the right, he has no reason to give up $\neg F_{n-1}$, or F_{n-2}, ... At first glance, it would therefore appear that the doxastic state is permissible.

At second glance, one might worry that perhaps Dom does have some reason to give up $\neg F_{n-1}$ upon learning F_n. For given that, say, the second stone fell, it is natural to ask what caused it to fall. And since one of the things that may have caused this is the first stone falling, perhaps Dom does have some reason to allow for the possibility that the first stone fell as well. This objection may be avoided, however, by modifying the example, at the cost of some additional complexity.

[12] The reason for using this notation is that it helps bring out more clearly the connection to counterfactual logic. This will help relating the present discussion to Fine's, and in particular means that his central proofs carry over to our setting without any changes.—The idea of interpreting a conditional in terms of belief revision in this way is again familiar from previous work, most notably in connection with the Ramsey Test; see e.g. Gärdenfors (1986, 1988, 1987) and Makinson (1990); see also (Fermé and Hansson, 2018: p. 85f).

The difficulty arises because in the case of the dominos, the truth of F_n would be *responsible* for the truth of F_{n+1}, and ultimately F_m whenever $m > n$. But this is an inessential feature of the example. Indeed, for roughly similar reasons, Fine has already described a version of the example which lacks this feature (2012a: p. 224f). Transposed to the belief revision setting, the case runs as follows. We imagine another doxastic agent, Matt. His relevant beliefs are that there is an infinity of matches m_1, m_2, ..., placed in causal isolation from one another, each of them in an environment maximally conducive to the match lighting upon being struck, but none of them actually struck. Now let S_n be the proposition that match m_n is struck, let L_n be the proposition that match m_n lights, and let W_n be the proposition that match m_n is wet. Let S be $S_1 \wedge S_2 \wedge \ldots$, so S says that each match is struck. Then F_n is $S \wedge ((W_n \wedge \neg L_n) \wedge (W_{n+1} \wedge \neg L_{n+1}) \wedge \ldots)$. So F_n says that each match is struck, but every match from n onwards is wet and does not light.

Note that for all n, F_n contains F_{n+1} as a conjunct. So in this version of the case, the dispositions ascribed in (D.+) are simply dispositions to believe conjuncts of conjunctive information received, and therefore clearly permissible. So let us turn to the dispositions ascribed in (D.−), and let us consider the instance $F_2 \Rightarrow \neg F_1$. Note first that since Matt believes each match to be in an environment maximally conducive to its lighting upon being struck, learning that the first match is struck (S_1) would give Matt good reason to believe that match 1 lights (L_1). So it seems rational for Matt to believe that L_1 upon learning that S_1. Now F_2 is the conjunction of S_1 with some information exclusively about other matches, believed by Matt to be causally isolated from match 1. None of this additional information seems in any way to undermine the support that S_1—match 1 is struck—offers for L_1—match 1 lights.[13] So it also seems rational for Matt to believe that L_1 upon learning that F_2. But now note that L_1 logically entails $\neg F_1$, since F_1 contains $\neg L_1$ as a conjunct. So since it seems clearly rational for Matt to form the belief that L_1 upon learning F_2, and L_1 logically entails $\neg F_1$, it also seems clearly rational for Matt to retain the belief that $\neg F_1$ upon learning that F_2. In other words, learning that F_2 not only provides no reason for Matt to give up the belief that

[13] Perhaps one might object that even though the matches are assumed to be causally isolated from one another, since according to F_n, the fates of matches n and onwards are so similar, it is still rational to suspect some kind of systematic explanation, which could then also suggest that earlier matches suffered the same fate. But it is not even necessary to suppose that events in the different regions are similar in this way. All we need is that S_n always says that some 'trigger'-event occurred, that L_n says that the corresponding standard result occurred, and that W_n says that some corresponding 'blocker'-condition obtained. (Cf. (Fine, 2012a: p. 225), see also (Fine, 2012b: p. 35, fn. 1).)

$\neg F_1$, it gives Matt additional support for that belief. Parallel considerations apply with equal force to the other instances of (D.−). I conclude that at least in this more complicated variant, the beliefs and dispositions to revise we have ascribed to Matt are jointly rationally permissible.

Now consider the infinite disjunction $F_1 \vee F_2 \vee \ldots$, and let us use F to abbreviate it. Assuming that the proposition that F is also a possible update, how should Matt be disposed to revise his beliefs upon learning that F? It is clear that there has to be *some* number n such that it is permissible for Matt to give up the belief that $\neg F_n$ upon learning that F. After all, if Matt were to retain each belief that $\neg F_n$ and add the belief that F, the resulting belief system would be inconsistent. We can also say something more specific, it seems to me. For it is hard to see how giving up $\neg F_n$ could be permissible for Matt for the case of, say, $n = 17$ but not for $n = 1$. So it also seems safe to assume that it is permissible for Matt to give up the belief that $\neg F_1$ upon learning that F.

We may summarize the central results of this section as follows: There is some permissible doxastic state which extends $\mathcal{D}$ and which, for some n, includes the disposition to give up the belief that $\neg F_n$ upon learning that F. In particular, there is some permissible doxastic state extending $\mathcal{D}$ and including the disposition to give up the belief that $\neg F_1$ upon learning that F.

In the next section, I will show that these results conflict with the AGM approach. Before that, let me address a kind of dismissive attitude towards these scenarios that some readers may be tempted to adopt. Clearly, both the domino- and the match-example are somewhat unrealistic. There are no infinite sequences of domino stones, and no infinite collections of matches in causal isolation from one another. So what, one might therefore ask, if our theory of belief revision has implausible implications with respect to such bizarre and silly cases? What matters, surely, is how belief systems relevantly similar to our own may be rationally revised, and the problematic kinds of doxastic states do not seem very similar to our own!

In response, it should be noted, firstly, that the specific subject matter of the above examples is of course not essential to the problem that they give rise to. All we need to generate that problem is an instance of the *general structure* exhibited by the cases of the dominos and the matches. So the objection can succeed only if all instances of this structure are silly. But that is not so. As Fine (2012b: p. 36) points out, one way to obtain more realistic instances is by considering, instead of infinite sequences of objects, infinite sequences of values of some quantity capable of continuous change, or at least taken by the agent to be so capable. Thus, we may consider an agent's beliefs concerning the flight of a missile believed to possess an automatic mechanism for correcting any deviations from its intended path (the example is Fine's). The propositions F_1, F_2, ... are now to the effect that the missile deviated by 1 inch off course, that the missile deviated by 1/2 inch off course, ... Since any deviation occurs in a continuous way, upon learning F_n the agent will believe F_m whenever $m \geq n$. But they may rationally retain the belief $\neg F_m$ whenever $m < n$, taking the mechanism to have prevented any greater deviation.

Another idea, more promising in our context of belief revision than in Fine's context of counterfactuals, is to construct an example using actual infinite sequences of abstract objects, such as the sequence of the natural numbers. What we would need is an example of a property with respect to which an ideal agent might initially believe that no number has it, and be disposed, upon learning that n has the property, to form the belief that m has it for all $m \geq n$, and to retain the belief that m does not have it for all $m < n$. Indeed, we might approximate the structure of the match example by letting F_n say that (a) for each n, attempts have been made to prove that n has the property, and (b) a proof has been found for each $m \geq n$. The supposition that this kind of situation could arise for some complicated number-theoretic property does not appear problematically unrealistic.

Secondly, the objection overestimates the role that infinity plays for the problem. As we shall shortly see, the relevant assumptions of the intensional approaches yield highly counter-intuitive results even in application to related, finitary contexts. Roughly speaking, the role of infinity is only to turn counter-intuitive results into contradictory ones. Relatedly, the approach I shall eventually propose deviates from its intensional rivals even in finitary contexts, and may be argued to be superior to them even on the basis of considering only finitary contexts.

4. Against the Possible Worlds Approach

We shall now show that the AGM approach is incompatible with the results of the previous section. To start, let us assume Matt is disposed to give up the belief that $\neg F_1$ upon learning that F:

(X_1) $F \not\Rightarrow \neg F_1$

This is rationally incompatible, given AGM, with

(1) $F_1 \Rightarrow F_2$
(2) $F_2 \Rightarrow \neg F_1$

To see this, note first that under the AGM approach, for any propositions P and Q, $P \Rightarrow Q$ holds iff $B * P$ entails Q, i.e. iff $B * P \subseteq Q$. So (X_1) implies that $B * F \nsubseteq \neg F_1$. By the definition of revision in terms of the plausibility ordering, $B * F$ comprises exactly the maximally plausible F-worlds, so (X_1) requires that some maximally plausible F-world be an F_1-world. Call that world w. By (1), every maximally plausible F_1-world is an F_2-world, so w is also an F_2-world. By (2), every maximally plausible F_2-world is a $\neg F_1$-world, so among the F_2-worlds, some world v must be more plausible than w. But every F_2-world is also an F-world, so v is a more plausible F-world than w, contrary to the assumption that w is a maximally plausible F-world. Since we found it to be rationally permissible for Matt to satisfy (X_1), this is a problem.

Moreover, by similar reasoning we can show that *any* instance of

(X_n) $F \not\Rightarrow \neg F_n$

is rationally incompatible, under the possible worlds approach, with (D.+) and (D.−), for there is no ordering of the worlds that satisfies the conditions (1)–(4) on plausibility orderings and that validates the dispositions in (D.+) and (D.−). In particular, any such ordering that respects (D.+) and (D.−) is such that there is no maximal F-world. For suppose w is an F-world, and let m be some number such that w is an F_m-world. Suppose for contradiction that w is a maximal F-world. Then in particular, w is a maximal F_m-world. By (D.+), every maximal F_m-world is also an F_{m+1}-world. By (D.−), no maximal F_{m+1}-world is an F_m-world. So w is not a maximal F_{m+1}-world, and hence not a maximal F-world after all.

Alternatively, as Fine shows (2012a: pp. 244ff), we can also derive all instances of $F \Rightarrow \neg F_n$ from (D.+) and (D.−) using only the following inferences rules, all of which are valid under the possible worlds approach:

Substitution	$P \Rightarrow Q \;/\; P' \Rightarrow Q$	[given that P and P' are logically equivalent]
Entailment	$/\; P \Rightarrow Q$	[given that P logically entails Q]
Transitivity	$P \Rightarrow Q,\; P \wedge Q \Rightarrow R \;/\; P \Rightarrow R$	
Conjunction	$P \Rightarrow Q,\; P \Rightarrow R \;/\; P \Rightarrow Q \wedge R$[14]	
Disjunction	$P \Rightarrow R,\; Q \Rightarrow R \;/\; P \vee Q \Rightarrow R$[15]	

[14] Fine also has an infinitary version of this rule, allowing us to infer $P \Rightarrow Q_1 \wedge Q_2 \wedge \ldots$ from $P \Rightarrow Q_1$, $P \Rightarrow Q_2, \ldots$. Using this rule we could show that conforming to the rules would lead, in the case at hand, to Matt's believing an outright contradiction upon learning F. But it seems bad enough if Matt ends up with an unsatisfiable belief system, accepting an infinite disjunction while rejecting each disjunct. For this result we only need the finitary rule.

The complete proof of this result is fairly long and complicated, so I shall refrain from reproducing it here. To see where the reasoning of the proof might best be resisted, and thus which rule might best be given up, it is more helpful to present it in more informal terms. And since the match-case is rather complex and hard to think about, in commenting on the various steps, I will use the dominos-example again. Let me first explain why giving up $\neg F_1$ in response to F would involve a violation of the rules. We may divide the reasoning into three main steps.

The first step is an application of *Substitution*, taking us from $F_2 \Rightarrow \neg F_1$—an instance of (D.–)—to $(F_1 \wedge F_2) \vee (\neg F_1 \wedge F_2) \Rightarrow \neg F_1$. In the domino case, this says that given that Dom is disposed to retain the belief that the first stone stands given the information that the second fell, he must also be disposed to retain that belief given the information that either the first and second, or not the first but the second stone fell.

The second step is to infer from $(F_1 \wedge F_2) \vee (\neg F_1 \wedge F_2) \Rightarrow \neg F_1$ that $F_1 \vee (\neg F_1 \wedge F_2) \Rightarrow \neg F_1$: Dom must also retain the belief that the first stone stands upon learning that the either the first stone fell or not the first but the second fell. The justification for this is that Dom is disposed to form the belief that F_2 given the information that F_1. Because of this, for Dom, learning F_1 and learning $F_1 \wedge F_2$ effectively come to the same thing, and the same is then true for learning $F_1 \vee (\neg F_1 \wedge F_2)$ and $(F_1 \wedge F_2) \vee (\neg F_1 \wedge F_2)$.

The third step is to infer from $F_1 \vee (\neg F_1 \wedge F_2) \Rightarrow \neg F_1$ that $F \Rightarrow \neg F_1$. Here, the idea may be described as follows. Learning F presents Dom with a choice: he needs to pick *some* stone s_n as the left-most stone for which to give up the belief that $\neg F_n$. Now $F_1 \vee (\neg F_1 \wedge F_2)$ says that either s_1 or s_2 is the first stone to fall. So learning $F_1 \vee (\neg F_1 \wedge F_2)$ presents Dom with a related choice: he needs to pick some stone $s_n \in \{s_1, s_2\}$ as the left-most stone for which to give up the belief that $\neg F_n$. Now the point is if Dom does not pick s_1 among the options s_1 and s_2, he cannot rationally pick s_1 among the options $s_1, s_2, \ldots$ Put another way, that $F_1 \vee (\neg F_1 \wedge F_2) \Rightarrow \neg F_1$ means Dom, prefers the scenario in which s_2 is the left-most stone to fall to the scenario in which s_1 is the left-most stone to fall. But giving up $\neg F_1$ in response to F would mean *not* preferring any alternative scenario to the scenario with even s_1 falling. So in particular, it would mean not preferring a scenario s_2 as the left-most stone to fall to the scenario with even s_1 falling.

---

15 As Fine points out, we actually require only a weaker rule with the added condition that P and Q be logically exclusive. The difference is not essential for present purposes, so for simplicity, I've here stated the stronger one.

So if Dom is to conform to the above rules, he must retain $\neg F_1$ upon learning F, and hence conclude that one of the other stones fell, i.e. $F_2 \vee F_3 \vee \ldots$ But the same considerations that prevent him from giving up $\neg F_1$ to accommodate F also prevent him from giving up F_2 to accommodate $F_2 \vee F_3 \vee \ldots$, so in the end he is prevented from giving up any $\neg F_n$. Resisting this final part of the argument seems hopeless. As mentioned before, it simply beggars belief that general constraints of rationality should prevent Dom, in the case at hand, from giving up $\neg F_1$ in response to F, while allowing him to give up, say, $\neg F_{17}$.

Applying AGM theory proper to the examples is not completely straightforward, since the examples involve infinite (conjunctions and) disjunctions, and AGM, strictly speaking, is concerned only with finitary propositional languages. Still, we may consider a trivial extension of AGM to languages with infinite conjunction and disjunction, in which we simply retain all the usual postulates. In this extension of AGM, the above rules can all be derived, and thus the proof that Dom and Matt won't be allowed to give up any belief of the form $\neg F_n$ can be carried out.

But we can also adjust the example so as to do without any infinitely long sentences. Instead, we may replace each infinite conjunction and each infinite disjunction used in our argument by a propositional letter, interpreted as expressing the same proposition as the infinitary sentence it replaces. If it is objected that these propositions may not be graspable by finite thinkers, we can instead let the propositional letters express the universal quantifications corresponding to the infinite conjunctions and the existential quantifications corresponding to the infinite disjunctions.

Since the dispositions ascribed in (D.+) and (D.−), under this modification, concern the same propositions as before – or perhaps quantificational counterparts – they are no less reasonable than before. So we still find that there can be no maximally plausible F-worlds. Since our background language now has a propositional letter true in exactly the F-worlds, it follows that no ordering of the worlds can satisfy ($\leq$1)–($\leq$3) together with the weakened version of ($\leq$4). And so we can infer by the mentioned equivalence that no AGM-revision operation can accord with (D.+) and (D.−) under the finitary replacement.

Most of the derivation given by Fine also still goes through under this modification. Infinitary sentences are involved only in the third step of the argument as described above, in which we infer $F \Rightarrow \neg F_1$ from

(3) $F_1 \vee (\neg F_1 \wedge F_2) \Rightarrow \neg F_1$

Formally, the way the derivation works is this. By *Entailment*, we also have

(4) $\neg F_1 \wedge \neg F_2 \wedge (F_3 \vee F_4 \vee \ldots) \Rightarrow \neg F_1$

By *Disjunction*, we obtain

(5) $(F_1 \vee (\neg F_1 \wedge F_2)) \vee (\neg F_1 \wedge \neg F_2 \wedge (F_3 \vee F_4 \vee \ldots)) \Rightarrow \neg F_1$

Now the point is that this big disjunction is logically equivalent to $F = F_1 \vee F_2 \vee \ldots$, so that by *Substitution* we may infer $F \Rightarrow \neg F_1$.

But now let F and F^3 be propositional letters expressing the proposition that some stone fell, and that some stone other than the first two fell, respectively. Then *Entailment* and *Disjunction* also give us

(4′) $\neg F_1 \wedge \neg F_2 \wedge F^3 \Rightarrow \neg F_1$
(5′) $(F_1 \vee (\neg F_1 \wedge F_2)) \vee (\neg F_1 \wedge \neg F_2 \wedge F^3) \Rightarrow \neg F_1$

Now the antecedent in (5′) is not logically equivalent to F, so we cannot infer $F \Rightarrow \neg F_1$ simply by an application of *Substitution*. But it is very plausible to assume that it must always be rationally permissible for the agent to treat $(F_1 \vee (\neg F_1 \wedge F_2)) \vee (\neg F_1 \wedge \neg F_2 \wedge F^3)$ as equivalent to F in his dispositions to revise. So we may simply make it a further non-logical *assumption* of the case, in addition to (D.+) and (D.−), that the agent's dispositions satisfy this condition. Given this assumption, we may then infer $F \Rightarrow \neg F_1$, and similarly for all other $\neg F_n$. In this way, even without the use of infinitary sentences, we obtain examples of ideally rational doxastic states that violate some of the AGM principles.

5. Against Intensionalist Responses

We saw that the doxastic states described, in virtue of satisfying (D.+) and (D.−), yield a violation of the condition ($\leq$4) of the possible worlds approach, requiring each set of possible worlds—or each expressible set of worlds in case of the weakened version—to have a maximally plausible member. One obvious idea for responding to the problem while retaining much of the original framework is therefore to drop this condition. There is even a precedent for this move for the case of counterfactuals, as the counterpart to ($\leq$4) in this setting is the so-called *limit assumption*, famously rejected by Lewis.

Without ($\leq$4), we can no longer define the revision by update P as the set of the maximally plausible P-worlds. How are we to define it instead? Lewis's proposal for truth-conditions for counterfactuals is of no help. Lewis takes $P \mathbin{\Box\!\!\rightarrow} Q$ to be true iff Q is true in all sufficiently close P-worlds, i.e. iff by restricting attention more and more to ever closer P-worlds, eventually we will be left only with Q-worlds. The simplest way to see this is to note that the Lewis-style truth-conditions for $P \Rightarrow Q$ actually validate all the above inference rules.[16]

A natural idea at this point is that the new belief state, in cases where there are no maximally plausible updates, should simply contain *all* the update-worlds.[17] At first glance, this may look attractive. It allows (D.+) and (D.−) to hold, while also allowing that the agent gives up $\neg F_1$ upon learning F, since some F-worlds are F_1-worlds. But at second glance it becomes clear that this suggestion throws out the baby with the bathwater. For the proposal does not allow our agent to have *any* beliefs, post-revision, save for those entailed by the update F. For example, our agent is not allowed to believe, post-revision, that if F_1 then F_2, since it is compatible with the truth of F that $F_1 \wedge \neg F_2$. But it is clearly rational in our scenario to retain the belief that F_2 if F_1, and so the proposal still misclassifies rational doxastic states as irrational.

Perhaps, then, we might give up on the idea that every rational revision function must be definable in terms of a plausibility ordering. Instead, we might say merely that any rational revision function must *conform*, in some suitable sense, to a plausibility ordering, and allow that there may be more than one revision function conforming to a given plausibility ordering. A natural first suggestion would be to take a revision function to conform to a plausibility ordering iff it maps any update P to the set of maximally plausible P-worlds if that set is non-empty, and to some upwards closed non-empty subset of P if not, where a subset P' of P is upwards closed iff P' includes every P-world that is more plausible than some world in P'.[18]

In terms of the inference rules employed in Fine's derivation, this proposal invalidates the *Disjunction* rule. In particular, it leads to the rejection of the inference from (3) and (4′) to (5′):

(3) $F_1 \vee (\neg F_1 \wedge F_2) \Rightarrow \neg F_1$

(4′) $\neg F_1 \wedge \neg F_2 \wedge F^3 \Rightarrow \neg F_1$

[16] There is a rule that is invalidated by adopting the Lewis-style truth-conditions, namely the infinitary version of the conjunction rule (cf. (Fine, 2012a: p. 225)). As mentioned before, this rule is not required for our purposes.

[17] This corresponds to the idea considered by (Fine, 2012a: p. 228f) of taking $P \Box\!\!\rightarrow Q$ to be true iff Q is true in all the closest *and all the stranded* P-worlds, where a P-world is stranded iff there is no closest world closer than it. Fine's most important objection against the proposal is analogous to my criticism in the main text.

[18] Probably, one should then impose some further constraints on how the choices of subsets for different updates have to relate. For instance, any world in the revision by $F_2 \vee F_3 \vee \ldots$ should probably also be included in the revision by $F_1 \vee F_2 \vee F_3 \vee \ldots$

(5′) $(F_1 \vee (\neg F_1 \wedge F_2)) \vee (\neg F_1 \wedge \neg F_2 \wedge F^3) \Rightarrow \neg F_1$

In terms of the AGM postulates, the proposal invalidates the postulate of *Superexpansion*, which says that the result of revising with a proposition P, conjoined with Q, entails the result of revising with $P \wedge Q$. To see how this fails, note that the result of revising with F, under the present proposal, is compatible with F_1, and remains so when conjoined with F_2. At the same time, since $F \wedge F_2$ is logically equivalent with F_2, the result of revising with $F \wedge F_2$ is not compatible with F_1, since the belief that $\neg F_1$ is retained in the revision by F_2.

Although an improvement over the previous attempts, this strategy is still unsatisfactory. For the proposal to be adequate, two conditions must be satisfied. Firstly, the complete extensions of the doxastic state that it classifies as permissible must really be so. Secondly, it must classify every permissible extension of the state as permissible. With respect to both conditions, there are good reasons to be skeptical.

Regarding the first condition, the problem is that the rejected applications of *Disjunction* and *Superexpansion* are intuitively very plausible. In the case of *Disjunction*, we assume that upon learning that $F_1 \vee (\neg F_1 \wedge F_2)$—all matches are struck, but all matches from the first or the second onwards are wet and do not light—Matt retains the belief that $\neg F_1$, and thus excludes the possibility that the first match is wet and does not light. He also retains that belief, obviously, upon learning that $\neg F_1 \wedge \neg F_2 \wedge F^3$. How can it then be rational for Matt not to retain the same belief—and thus to allow for the possibility that the first match is wet—upon learning the disjunction of these two pieces of information?

The case of *Superexpansion* seems even more compelling. We take for granted that Matt gives up the belief that $\neg F_1$—and so allows for the possibility that the first match is wet and does not light—upon learning that all matches are struck, but all matches from some match onwards are wet and do not light. But then how can it be rational to retain the belief that $\neg F_1$—and thus exclude the possibility that the first match is wet and does not light—upon receiving the same information, with the addition that either match 1 or match 2 is the first match to be wet and fail to light?

Regarding the second condition, there are strong reasons to think that there are other permissible extensions of the doxastic state than those envisaged under the present proposal. For instance, it seems very plausible that it should be permissible for Matt's doxastic state to be such that

(6) $F_1 \vee \ldots \vee F_{100} \not\Rightarrow \neg F_{99}$

That is, it should be permissible for Matt to be disposed to give up the belief that $\neg F_{99}$ upon learning that $F_1 \vee \ldots \vee F_{100}$. For consider what $F_1 \vee \ldots \vee F_{100}$ says. It says that all matches are struck, and that for some match m_k among the first 100, all matches from m_k onwards are wet and do not light. It would seem quite bizarre for Matt, upon receiving this information, to retain the belief that $\neg F_{99}$, and thus to conclude that m_k must have been m_{100}, i.e. that it must have been match 100 that is the first in the sequence to be wet and fail to light. It certainly does not seem as though having the dispositions in (D.+) and (D.−) *requires* Matt to respond in this way to the information that $F_1 \vee \ldots \vee F_{100}$.[19]

Similarly, it seems that it should be permissible for Matt to be such that

(7) $F_1 \vee F_2 \not\Rightarrow \neg F_1$

That is, it should be permissible for Matt to be disposed to give up the belief that $\neg F_1$ upon learning that all matches are struck, and either all the matches, or all matches from the second onwards, are wet and do not light.

These intuitions are in conflict with the principle of intensionality. For under the interpretation given in the match case, F_1 contains F_2 as a conjunct, and so $F_1 \vee F_2$ is logically equivalent to F_2—and upon learning that F_2, by (D.−), Matt is disposed to retain the belief that $\neg F_1$. Likewise, all of $F_1, \ldots, F_{99}$ contain F_{100} as a conjunct, so $F_1 \vee \ldots \vee F_{100}$ is logically equivalent to F_{100}—and upon learning that F_{100}, by (D.−), Matt is disposed to retain the belief that $\neg F_{99}$. Let us see, then, where we can get by dropping the assumption of intensionality and trying to accommodate these intuitions.

6. Towards a Hyperintensional Solution

We begin by sketching a general method for revising one's beliefs that Matt might be seen to follow and that would lead to his conforming to the intuitions just observed. Both (6) and (7) concern how Matt revises his beliefs by a disjunctive piece of information. A very natural idea is that he does this by forming the disjunction of the results of revising his beliefs by each disjunct. Thus, if we write B for Matt's initial beliefs and $*$ for his revision function, the idea is that $B * (F_1 \vee \ldots \vee F_{100}) = (B * F_1) \vee \ldots (B * F_{100})$, and $B * (F_1 \vee F_2) = (B * F_1) \vee (B * F_2)$.[20] If so, since $B * F_{99}$, for example, does not entail

[19] Note that this problem arises in exactly the same way in a finitary version of the example, as the assumption that there are infinitely many matches does no work here. This is way I said at the end of the previous section that the role of infinity is merely to turn counter-intuitive results—like this one—into contradictory ones.

[20] This assumes that belief systems are among the kinds of things to which the operation of disjunction can be applied. This is unproblematic if belief systems are identified with propositions, and slightly less straightforward when belief systems are identified with sets of sentences, though it is clear enough

that $\neg F_{99}$, then neither does $B * (F_1 \vee \ldots \vee F_{100})$, in line with (6). And likewise since $B * F_1$ does not entail that $\neg F_1$, then neither does $B * (F_1 \vee F_2)$, in line with (7).

Borrowing a term from Fine (2012b: p. 52), we may call this the method of *wayward revision*, since it involves revising, one by one, by each disjunct of the update, i.e. by each way for the update proposition to be true. (And here, as in Fine, waywardness is considered a good thing.) Revising in this way means that every disjunct of the update is accommodated by the agent in the sense that there is some way for the revised belief system to be true under which that disjunct of the update is true. In other words, for each disjunct Q of the update P, according to the wayward revision by P, it might be that Q. Now to adopt the view that it might be that Q on some occasion for revision—even if one's belief previously excluded the possibility that Q—is to treat the occasion as telling one that it might be that Q. In this sense, the method of wayward revision seems to depend on a principle about updates that we may roughly express like this:

(M) A situation with update $P \vee Q$ is a situation telling the agent that it *might* be that P, and that it *might* be that Q.

Whenever a pair of situation and update satisfy (M) with respect to all disjuncts of the update, I shall say that the update is *mighty* in that situation. The fully general claim (M) is then that updates are always mighty. A central feature of the approach to belief revision that I want to propose is that it endorses principle (M).

Why should one endorse that principle? One consideration in favour of (M)—not the only one—is that it makes sense of our above described intuitions: our intuitive verdicts regarding the rational ways to revise by updates such as $F_1 \vee \ldots \vee F_{100}$ or $F_1 \vee F_2$ in our puzzle cases seem to arise from a tacit assumption that the updates are mighty in the situations under considerations.

Now it might be objected that it is a mistake to let oneself be guided by these intuitions, since they are simply owed to certain *pragmatic* effects. The thought might be spelled out as follows: The update is supposed to capture the total information received by the agent in the relevant situation. To say that an agent receives the information that $P \vee Q$ pragmatically conveys that, in the situation in question, the agent is given some reason to allow for the possibility that Q. For suppose the agent is given no such reason. Then it will normally be wrong to say that the total information received is that $P \vee Q$, since the agent will then also have received the information that P, which is normally stronger than the information that $P \vee Q$. The exception is if, as in our examples, the

how the notion of disjunction should be extended from sentences to sets of sentences. Still, we shall always be thinking of belief systems as propositions.

propositions that P and that $P \vee Q$ are logically equivalent, since Q is of the form $P \wedge R$. But in such cases it will still be highly misleading to say that the information received is that $P \vee Q$, since it is hard to see what the point could be of presenting the information in this disjunctive form except to convey that the agent is given some reason to allow for the possibility that Q. Still, that the agent is given such a reason is *merely* pragmatically conveyed by the statement that the total information they received is that $P \vee Q$. It is not, or so the objection goes, part of the semantic content of that statement.

The objection misses the point. For all I want to argue here, it may well be that as a sentence of ordinary English, an instance of 'the total information the agent received is that P or Q' does not semantically imply that the agent is given reason to allow that it might be that P, and that it might be that Q. But our ultimate goal here is not to analyse ordinary discourse about people receiving information, it is to develop an adequate theory of belief revision, i.e. to adequately capture the general rationality constraints on doxastic states. As part of this, we require some means to pair occasions for revision with propositions—which we call the updates—in such a way that only dynamically equivalent occasions are assigned the same proposition. A rough and ready informal characterization of a suitable pairing uses talk of what the agent learns, or what information they receive. But in developing our theory of belief revision, we may have occasion to clarify or refine that rough characterization in certain ways. How this should be done depends more on the theoretical requirements of a theory of belief revision, and less on the available readings of the relevant locutions in ordinary English.

What I wish to claim is, firstly, that we *can* pair occasions for revision with propositions as their updates in such a way that (i) only dynamically equivalent situations are paired with the same update, and (ii) updates are always mighty. Secondly, I claim that for the purposes of theorizing about rational belief revision, it is *beneficial* to characterize occasions for revision in terms of these mighty updates. The distinction between pragmatic and semantic implications has little bearing on these claims. In defence of these claims, I will develop a conception of updates as mighty which is based on the framework of truthmaker semantics (§7), formally characterize a class of permissible doxastic states within the truthmaker framework and show that they satisfy versions of all the usual AGM postulates save for intensionality (§8), and finally highlight what I take to be the important general advantages, apart from our puzzle cases, of the resulting approach and especially the conception of updates as mighty (§9).[21]

[21] [Footnote removed for blind review]

7. Mighty Truthmaker Updates

To begin, let me make two important initial clarifications regarding the notion of the update which are independent of any issues around mightiness or hyperintensionality. The first is that I take the update to represent the information the agent *takes* themselves to obtain in the given situation, or perhaps better: the information the agent *treats* the situation as providing them with. In particular, if there is also a distinct notion of what information a situation *really* provides a given agent with, whether or not the agent regards and treats the situation accordingly, then that is not what I intend to capture in the update. An example may help to make this clearer. Suppose I have the kind of visual experience that would normally lead to me coming to know that my neighbour is walking towards my house. The experience is caused in the appropriate sort of way by my neighbour walking towards my house, my visual system is as it should be, and so on. But suppose further that I have misleading evidence to take my visual system to be compromised, and thereby to doubt the veridicality of my experience. In one sense, perhaps, this is a situation in which I receive the information that my neighbour is walking towards my house—it is just that that circumstances are such as to (rationally) prevent my uptake of that information. But in the sense I intend, this is not a situation in which I receive the information that my neighbour is walking towards my house. For it is not a situation which I treat as giving me this information. Conversely, a situation in which someone tells me that *P*, and I trust the speaker, would be a situation in which, in the intended sense, I receive the information that *P*, even if the speaker is actually lying, and it is false that *P*.

A second, in some ways complementary clarification is that I take the update to represent what the subject treats *the situation*—on its own, as it were—as telling them. Consider a version of the previous scenario in which I have no doubts about my visual system and accordingly come to the belief that my neighbour is walking towards my house. Suppose further that my wife previously told me that my neighbour is away on holiday, leading me to conclude that my wife was mistaken. In one sense, perhaps, I might be said to treat the situation as providing me with the information that my wife was mistaken. But this seems to be a case in which, *in the course of revising* my belief system in the light of the new information, I come to acquire this belief. It is not a case in which the relevant belief is part of the information I treat the situation *on its own* as providing me with.

Both these stipulations are reasonable independently of the questions of mightiness and hyperintensionality. Unless we make the first stipulation, it is doubtful that rationality requires the agent to come to believe the update.[22] Unless we make the second stipulation, we lose the distinction between the interpretation of a situation by an agent on the one hand and the resulting adjustment of their previous beliefs on the other.

We may thus think of the process of belief revision as divided into two stages. The first stage consists of the agent *interpreting* the situation in which they find themselves, and deciding what to take it as telling them. The second consists of the agent revising their beliefs in light of what they've taken the situation to tell them. The role of the update is to represent the outcome of stage one. In explaining our conception of the update, what we need to explain is therefore what it says about how the agent interprets the given situation that we are assigning to it a particular update.

The conception of updates I wish to propose is intended to render them mighty, so that by assigning to a situation the update $P \vee Q$, we are saying, among other things, that the agent interprets the situation as telling them that it might be that P, and that it might be that Q. The condition of the situation telling the agent that it might be that P here is to be understood in a specific, comparatively demanding way. In a weak sense, we might say that the situation tells the agent that it might be that P whenever the situation, as interpreted by the agent, does not—actively and by itself, as it were—exclude the possibility that P. A more natural interpretation of the condition is more demanding. It requires, we might say, that the situation explicitly presents it as a possibility that P, that it being the case that P would (at least) help account for the situation, or that it being the case that P would (at least) partially constitute the truth of what the agent takes the situation to tell them.

The distinction is difficult to define in independent, non-metaphorical terms, but it is clear and familiar enough. An example may help to illustrate the idea. Suppose my neighbour has twin sons, Bob and Bill. Suppose further that I see someone walking towards my house, and that I see them well enough to be able to tell that it is definitely either Bob or Bill, but I can't tell which. So I take the situation to tell me, among other

[22] This requirement is implicit in the rule of *Entailment*, and captured in the AGM postulate *Success*. See Stalnaker (2009) for a similar approach to justifying the *Success* postulate. (I do not mean here to exclude the possibility of fruitfully theorizing about belief revision on the basis of a different conception of the update, not subject to the requirement that the agent takes themselves to come to know the update. But this would constitute a more radical departure from the AGM tradition than I wish here to consider. In the literature, approaches of this sort often go under the label of *non-prioritized* belief revision; for a brief introduction see (Hansson, 2017: § 6.3).)

things, that Bob or Bill is coming over. Consequently, some propositions are *incompatible* with the situation as I interpret it, such as any state to the effect that both Bob and Bill are away on holidays, say. Some propositions are *merely compatible* with the situation as I interpret it, such as the proposition that it is sunny in Ohio. And some propositions are explicitly presented as possiblities by the situation, such as the proposition that Bob is coming over, and the proposition that Bill is coming over. These are propositions we might describe as (partially) accounting for the situation I find myself in, as I interpret it, as propositions whose truth would partially constitute the truth of what I take the situation to tell me. Let us call propositions in this final category *explicit possibilities* of the situation (under the agent's interpretation[23]), and those in the former category merely *implicit* possibilities.[24]

The distinction between explicit and implicit possibilities is relevant to how an agent may rationally revise their beliefs. If a proposition is an explicit possibility in a situation, then the situation provides some *reason* for the agent to allow for the possibility of the proposition's being true, even if their original belief system excludes that possibility. Thus, in the example, even if I initially believed both Bob and Bill to be away on holiday, the situation provides some reason for me to allow for the possibility that Bob is coming over, and it provides some reason for me to allow for the possibility that Bill is coming over. But if a proposition is a merely implicit possibility, then the situation does not give the agent reason to allow for the possibility that it is true. If in our example I originally believed it not to be sunny in Ohio, then the situation provides no grounds whatsoever to subsequently allow for the possibility of it being sunny in Ohio.

Crucially, the condition that the situation tells the agent that it might be that P in (M) is to be understood as requiring that the proposition that P is an *explicit* possibility in the

23 This qualification will henceforth usually remain tacit.

24 The distinction between what I have called explicit and implicit possibilities in a situation may be compared to the distinction between the strong and weak permissions of a system of norms (cf. (von Wright, 1963: p. 90)), where an action is weakly permitted iff it is compatible with the system of norms, and strongly permitted iff it is actively singled out, as it were, as permitted by the system of norms. The difficulties in capturing these distinctions within an intensional framework are likewise parallel. Fine (2018) proposes a truthmaker semantics for statements of permission that is sensitive to the distinction, and captures it in much the same way that I propose below. In §8 of that paper, Fine also adresses the problem of deontic updating and notes the connection to belief revision. The approach to deontic updating Fine sketches is related to the approach to belief revision to be described below, but with a simple mereological construction taking the place of the transition relation invoked below. Related approaches to deontic updating are pursued in Yablo (2011) and Rothschild and Yablo (202x), who also draw the connection to belief revision. [Rest of footnote removed for blind review.]

situation. So under a conception of updates as mighty, to say that the update in a given situation is $P \vee Q$ is to say, among other things, that the agent is given some reason, in that situation, to allow for the possibility that P, and to allow for the possibility that Q.

We can now argue that if updates are mighty, they *must* be individuated in a hyperintensional way. For assuming intensionality, any given update P can also be written as $P \vee (P \wedge Q)$, for arbitrary Q. Assuming mightiness, it follows that in any situation with update P, the agent is told that it might be that $P \wedge Q$, and hence that it might be that Q, for arbitrary Q. Whatever P is, there will be few if any such situations. Conversely, it seems most situations will not be representable by a mighty intensional update. It seems safe to conclude, therefore, that a conception of updates as mighty requires a hyperintensional way of individuating updates, in particular one that allows us to distinguish between pairs of the form P and $P \vee (P \wedge Q)$.[25]

I propose that we model updates as propositions as conceived within the framework of truthmaker semantics.[26] Within this theory, propositions are characterized not (merely) in terms of the possible worlds at which they are true, but in terms of the possible *states* which *make* them true.[27] Informally, a possible state may be thought of as a (proper or improper) *part* or *fragment* of a possible world, but officially the notion is a primitive of the theory. States are taken to be ordered by part-whole ($\sqsubseteq$), and some states s_1, s_2,

[25] Since AGM is based on an intensional conception of the update, it would seem to follow from this that AGM updates cannot be considered mighty. On the other hand, one might argue that the AGM method of revision does reflect a conception of updates as mighty. The reasoning is this. As will become clearer in the next section, regarding an update as mighty means that for each disjunct P of the update, *absent special reasons to the contrary*, P must be accommodated as a possibility. Now under the AGM account, the agent must accommodate a disjunct P unless they consider no P-worlds to be among the most plausible update-worlds. To the extent that this is a special reason not to accommodate P, AGM revision embodies a view of the update as mighty. Indeed, one might think this is exactly what goes wrong in our puzzle cases. AGM lets us retain $\neg F_1$ upon revising by F_2 only if we have special reasons to discard the F_1-worlds among the F_2-worlds. So in this way, the update F_2 is treated as mighty, and as identical to $(F_1 \wedge F_2) \vee F_2$. But we can be in a situation where it is fine just by default to accept F_2 and retain $\neg F_1$, because F_1 is merely compatible with the update F_2, and not an explicit possibility.

[26] A semantics of this sort was first formulated in van Fraassen (1969). In recent years, the approach and its various applications have been further developed by Fine and others. Fine's 2017a; 2017b offer the best general presentation of the theory. The following brief introduction is indebted to these works.

[27] A formally precise presentation of the framework is given in appendix A. For many applications of truthmaker semantics—including, I believe, some applications related to belief revision—, it is useful also to allow for a multiplicity of impossible states. For our present concerns, however, impossible states are not essential, though it will be convenient to assume that there is a single impossible state.

... are said to be *compatible* if there is a possible state that contains all of them as parts. It is assumed that there is always a smallest state to contain some given states s_1, s_2, ..., which we call their *fusion* $\bigsqcup\{s_1, s_2, \ldots\} = s_1 \sqcup s_2 \sqcup \ldots$[28] We may recover a notion of a possible world as the notion of a maximal possible state, i.e. a possible state that contains every state it is compatible with.

An exact truthmaker of a proposition is a state that is not only modally sufficient for the truth of the proposition, but also *responsible* for it. Thus, the state of it being sunny in New York is not an exact truthmaker of the proposition that 2+2=4. In addition, to be an exact truthmaker of a proposition, a state must be *wholly relevant* to the truth of the proposition. Thus, the state of it being sunny and cold in New York is not an exact truthmaker of the proposition that it is sunny in New York, since it contains an irrelevant part—the state of it being cold in New York—and therefore fails to be wholly relevant. The condition of being wholly relevant renders exact truthmaking non-monotonic: a given state may exactly verify, i.e. be an exact truthmaker of, a given proposition, without some bigger state also exactly verifying the same proposition. Relatedly, if a state s is an exact truthmaker of some proposition, then we may conclude that the proposition is in some good sense *about* the whole state s (though not in general *only* about s).[29]

This understanding of truthmaking suggests a particular account of the truthmakers of disjunctions and conjunction: A state makes a disjunction true iff it makes one of the disjuncts true, and it makes a conjunction true iff it is the fusion of truthmakers of the conjuncts.[30,31] Under this account, we can make the required distinction between P and $P \vee (P \wedge Q)$. Any fusion of a truthmaker of P and a truthmaker of Q is truthmaker of

28 When s_1, s_2, ... are incompatible, this will be the impossible state.

29 For much more on the relation of (non-monotonic) truthmaking to the notion of aboutness or *subject matter*, see Yablo (2015); Fine (2017b, ms).

30 There is also an alternative, *inclusive* clause for disjunction, in which the fusion of truthmakers of each disjunct is also considered a truthmaker. In some applications of truthmaker semantics it is preferable to work with the inclusive conception of disjunction, but as we shall see shortly, for the present application there are specific reasons not to do so.

31 Readers may wonder about the case of negation. The simplest approach is to associate any given proposition with both a set of exact truthmakers, and a set of exact falsitymakers, and to let negation ‘flip’ the two sets. For now, since none of the AGM postulates involves negation, we may set negation to one side. (Negation does of course play an important role in the relation between AGM-style revision and another important AGM-operation, namely contraction, which corresponds to the mere removal of a belief. There are important questions about the treatment of contraction and similar operations under a truthmaker approach, as well as about the matter of negation, but discussion of these will have to wait for another occasion.)

$P \vee (P \wedge Q)$, but since a truthmaker of Q will not in general be relevant to the truth of P, such a fusion will not in general be a truthmaker of P. In particular, we can distinguish between, for example, the logically equivalent F_2 and $F_1 \vee F_2$ in the match example. For by the clause for disjunction, every exact truthmaker of F_1 will be a truthmaker of $F_1 \vee F_2$. But since F_1 by definition has $W_1 \wedge \neg L_1$ as a conjunct, by the clause for conjunction, any such truthmaker will contain a part that makes true $W_1 \wedge \neg L_1$, the proposition that the first match is wet and does not light. That state will be irrelevant to the truth of F_2, and therefore no truthmaker of F_1 will be an exact truthmaker of F_2.[32]

We can now say which truthmaker proposition we take to be the update on a given occasion for revision. First, note that a division between explicit and implicit possibilities can also be made at the level of states. A state is an (at least) implicit possibility if it is compatible with what the agent takes the situation to tell them, and it is an explicit possibility if it also partially constitutes the truth of, i.e. partially *makes* true, what the agent takes the situation to tell them.[33] Among the explicit possibilities, we may then further distinguish between those that *merely partially* make true what the situation tells

[32] The central feature of the truthmaker framework is thus its use of a concept of *relevant* truthmaking, which makes it possible to capture various relationships of relevance between propositions. Relatedly, the distinctive features of the truthmaker-based approach to belief revision developed here can also be put in terms of relevance. On the conception of the update as mighty, the update $P \vee (P \wedge Q)$ is relevant to a prior belief in $\neg Q$, whereas the corresponding update P need not be so relevant. In particular, in our example, the update $F_2 \vee (F_1 \wedge F_2)$, but not the update F_2, is regarded as relevant to the belief that $\neg F_1$, and so the agent is permitted to be disposed to give up that belief in processing the former update while not being so disposed with regard to the latter update. The claim that AGM is not appropriately sensitive to the matter of which existing beliefs a given update is relevant to has also been made by earlier authors; most notably by Parikh (1999), whose proposal for extending AGM by a relevance axiom has been the subject of extensive discussion and refinements, cf. e.g. Kourousias and Makinson (2007); Makinson (2009). A proper comparison of the present approach with this tradition or other 'relevantist' criticisms of AGM is beyond the scope of this paper, but it may be worth mentioning two significant points of difference. Most of the work in the tradition initiated by Parikh embraces intensionality and accordingly does not adopt a conception of the update as mighty. That tradition also tends to follow a syntactically driven approach to understanding relevance (an exception is Peppas et al. (2004), providing a system-of-spheres semantics for Parikh's relevance axiom), whereas the present approach is chiefly driven by semantic concepts and considerations. It would be very interesting to study the relation between these approaches more deeply. One might try, for example, to formulate a suitably hyperintensional version of the relevance axiom and investigate whether it may be satisfied under some version of the present approach.

[33] Note that 'partial' hear means *part of* rather than *has as part*. Thus, by a partial truthmaker I mean something which is *part of* a truthmaker rather than something which has a truthmaker as a part.

the agent, and those that *fully* make true what the situation tells the agent. In our example, what I take the situation to tell me is perhaps not exhausted by the claim that either Bob or Bill are coming over. Perhaps I also see that Bob or Bill—whoever it happens to be—is wearing a black sweater and a red hat. Let us call explicit possibilities that fully make true what the situation tells the agent *complete*, and the others *incomplete*. The *truthmaker-update*—short: tm-update—in a given situation, as interpreted by the agent, is then the set of all and only the situation's complete explicit possibilities.[34]

Since the tm-update includes *only* explicit possibilities, a situation with tm-update P tells the agent, for each state $s \in P$, that s *might* obtain, and hence for each disjunct Q of P, that it *might* be the case that Q. Since the tm-update comprises *every* complete explicit possibility, moreover, a situation with update P tells the agent that it *must* be the case that P: the situation is taken by the agent to rule out any scenario in which it is not the case that P. We may summarize the point by saying that tm-updates are both *musty* and *mighty*.

By way of comparison, consider how an intensional conception of the update might be obtained. The obvious answer would seem to be as follows. Given an agent's interpretation of a situation, we divide the possible worlds into two exclusive and exhaustive categories. To the first belong those worlds that are compatible with the situation, under the agent's interpretation, and to the second belong the others. The *possible worlds update*—short: pw-update—is the set of the former worlds. Then pw-updates are certainly also musty: given that every world that is compatible with the situation is included in the update, we can conclude that the situation tells the agent that one of the update-worlds must obtain. But in contrast to tm-updates, which are musty and mighty, pw-updates are merely musty. For as we saw above, intensional updates cannot be mighty in the demanding sense in which tm-updates are.

Note that under our conception of tm-updates, assuming as given two situations with logically equivalent tm-updates P and Q that differ with respect to their truthmakers, there is nothing mysterious about why these situations can be dynamically inequivalent even assuming the agent knows the updates to be logically equivalent. That the agent

[34] Note that this set is plausibly not closed under fusion. For instance, Ben's coming over and Bob's coming over may each be explicit possibilities without Ben and Bob both coming being one. That is the reason why I think that in the application to belief revision, we need to allow for truthmaker propositions that fail to be closed under fusion, and relatedly to opt for the non-inclusive clause for disjunction, on which fusions of verifiers of the disjuncts are not automatically verifiers of the disjunction; cf. footnote 30 above.

knows that P and Q are logically equivalent means they know that it is absolutely impossible for P to be true without Q being true as well, and vice versa. A situation with tm-update P is one in which the agent takes themselves to learn that P. Knowing P to be equivalent to Q, they will also conclude that Q. Similarly in a situation with update Q. But *how* the belief that P, or the belief that Q, may appropriately be incorporated in these situations depends also on what the situations tell the agent about what *might* be the case. Given that P and Q have different truthmakers, situations with tm-updates P and Q will differ in this regard, and may therefore differ with respect to their range of rational responses.[35]

8. Revision

Given the proposed conception of updates as sets of truthmakers, how can we characterize the rationally permissible ways to revise a belief system by an update? First, we need to decide how to model belief systems within our revised setting. Although the issue calls for extended discussion, for present purposes we may adopt a policy of keeping this as simple as possible, and of minimizing deviation from the AGM approach, so that we may see how much, or how little, of that approach we are forced to give up to accommodate the problem cases. We shall therefore continue to model a belief state by the set of possible worlds at which it is true. Thus, the update will be the only source of hyperintensionality under the resulting approach.[36]

In imposing rationality constraints on doxastic states, we follow a similar strategy as the possible worlds approach in that we demand that the revision function be definable in a certain way. We suggested above that Matt might plausibly be seen to revise by disjunctions by disjoining revisions by the disjuncts. Within the truthmaker framework, a disjunct of an update is any subset of the update, and the disjuncts of the update which are not themselves disjunctive are the subsets with exactly one truthmaker as member.

35 We might compare the situation to the one in approaches to revision using belief bases, which are sets of sentences not (normally) closed under logical consequence. There, a distinction is made between, roughly speaking, sentences an agent believes to be true purely because they follow logically from other sentences the agent believes and sentences an agent believes to be true on (partly) independent grounds. The view is that rational revision is sensitive to this difference, and different but logically equivalent belief bases may rationally be revised differently. Just as in our case, the view is fully compatible with a view of agents as logically omniscient. See e.g. (Hansson, 1999: pp. 17ff).

36 That being said, my own view is that an ultimately more satisfactory approach may be obtained by also embracing hyperintensionality with respect to the belief system and representing an agent's beliefs by their exact truthmakers rather than all the verifying worlds. [Removed for blind review.]

So the suggestion is, in effect, to take the revision by an update to be the disjunction of the revisions by the individual truthmakers of the update. In this way, we obtain what we called the wayward revision of a belief system by an update.

Under certain circumstances, however, it may be rationally permissible for an agent to deviate from the method of wayward revision. The idea is that one may take a situation to tell one that it might be that P, and at the same reasonably hold that one knows better, as it were—that information one possesses independently of the given occasion for revision, and that is not undermined by the new information obtained, may justify continuing to exclude the possibility that P, even if the situation on its own is taken to explicitly present P as a possibility. First of all, one might so interpret a situation as to assign it the update $P \vee Q$, where P but not Q is compatible with one's previous beliefs. In the Bob-and-Bill case, for example, I might take the situation to tell me that Bob might be coming (P) and Bill might be coming (Q), when my beliefs are compatible with the former but not the latter possibility. In such a case, it is permissible for me to disregard the revision by Q and simply select as my new belief system the revision by P.[37] Second of all, even if all disjuncts of the update are incompatible with the agent's current beliefs, those beliefs may exclude the revisions by some disjuncts much more firmly than others, and it may then be rational for the agent to disregard the latter. In the context of the dominos, a plausible example might be the update to the effect that either all the stones fell, or exactly the odd-numbered stones fell. Given the setup of the case, any world verifying the second disjunct might seem a so much more remote possibility than worlds verifying the first disjunct that it may justifiably be disregarded. This suggests a modification of the simple method of revision, whereby revisions by disjunctive updates are constructed by first forming disjunction of the revisions by each disjunct, and then applying a "plausibility filter", discarding those disjuncts that are regarded sufficiently less plausible than others. Just like we did under the possible worlds approach, therefore, we may appeal to a plausibility ordering of the worlds, and let $B * P$ comprise only the most plausible worlds in the wayward revision of B by P.

It needs to be emphasized, however, that while from a formal perspective the plausibility orderings used here are just like those used in AGM, their representational role is quite different, and much less central to the overall account. In particular, under the

[37] Indeed, it is standardly assumed that this is not only permissible but mandatory. Specifically, the AGM postulate of *Vacuity* demands that no beliefs be given up in incorporating information compatible with the agent's current beliefs.

present approach, an agent may consider two initially excluded worlds equally plausible and yet, after a rational revision, continue to exclude one of them, while no longer excluding the other. Indeed, as will become clearer below, this is exactly what allows us to deal in an intuitively satisfactory way with the puzzle cases.

It remains to characterize the rationally acceptable ways to revise a belief state by a single truthmaker. A natural idea is to once more take a leaf out of Fine's semantics for counterfactuals (cf. 2012a: pp. 236ff), and to postulate a *transition* relation that encodes, roughly speaking, how each of the various worlds in the belief state may be adjusted upon revision by any given input state.[38] We write $s \rightarrow_b w$ to say that world w is a revision of world b by state s, and define the *wayward* revision $B \circ P$ of belief state B by update P as $\{w : p \rightarrow_b w$ for some $p \in P$ and $b \in B\}$. The final revision is then obtained by applying the plausibility filter. Where X is a set of worlds, we let $g(X)$ be the set of the maximally plausible members of X, and define $B * P$ as $g(B \circ P)$.[39]

Note that the revision operations of the usual possible worlds account constitute a special case of our revision operations, which corresponds to the condition that a world w is a revision of another world b by a consistent state s iff w contains s as part. Then the wayward revision of any belief state is simply the set of worlds at which the update is true, and the final, filtered revision is the set of the maximally plausible update-worlds. Thus, the way our present account improves on the possible worlds account, is by allowing the transition relation to narrow our focus from the start on some subset of the update-worlds, and to do so in a way sensitive to the exact truthmakers of the update.

38 Although in its use of a transition relation, the present approach thus maintains a strong parallel to Fine's semantics for counterfactuals, it should be noted that there is no counterpart in the latter to our use of plausibility orderings. Roughly speaking, while I propose to divide the work done by plausibility orderings under the possible worlds approach between plausibility orderings and a transition relation, Fine proposes to let transition do all the work of the similarity ordering in the possible worlds analysis of counterfactuals. I suspect that by using a similarity ordering in the account of counterfactuals, much as we use a plausibility ordering here, we might be able to avoid the difficulties for Fine's semantics raised in Embry (2014).

39 It is worth mentioning here that by appealing to the mereological as well as modal profile of states, it is possible to define in logical terms certain *defaults* for transition and plausibility. For instance, we might say that by default, a world w transitions to another v upon revision by a state s iff v is maximal among the s-containing worlds with respect to its mereological overlap with w. For plausibility, a natural default is to take all worlds incompatible with the current belief system to be equally plausible. In this way, the truthmaker approach allows us to give a purely logical characterization of a non-trivial operation of belief revision. [removed for blind review]

To illustrate the idea, we sketch a truthmaker model of Dom's doxastic state in the dominos example.[40] For simplicity, we let our worlds be built up purely from states of the form f_n—stone n falls—and $\overline{f_n}$—stone n does not fall. Then Dom's initial belief has just one verifier, the state $b = \bigsqcup\{\overline{f_n} : n \in N\}$. Its revision by the proposition that F_2, with its sole verifier f_2, will comprise exactly the maximally plausible worlds w with $f_2 \rightarrow_b w$. Its revision by the proposition that $(F_1 \wedge F_2) \vee F_2$, with is two verifiers f_2 and $f_1 \sqcup f_2$, will comprise exactly the maximally plausible worlds w with either $f_2 \rightarrow_b w$ or $f_1 \sqcup f_2 \rightarrow_b w$. We capture the fact that Dom takes the falling of any stone to lead to the falling of every subsequent stone by letting $f_1 \sqcup f_2 \rightarrow_b w$ hold if and only if $w = \bigsqcup\{f_n : n \in N\}$. The fact that Dom takes the falling of any stone not to support the falling of any previous stone is captured by letting $f_2 \rightarrow_b w$ hold if and only if $w = \overline{f_1} \sqcup \bigsqcup\{f_n : n \geq 2\}$. More generally, we capture Dom's dispositions to respond to a proposition of the form F_n by letting $f_n \rightarrow_b w$ hold iff $w = \bigsqcup\{\overline{f_m} : m < n\} \sqcup \bigsqcup\{f_m : m \geq n\}$. To accommmodate the fact that Dom is disposed to make room for the possibility that F_1 upon learning that F, or learning that $(F_1 \wedge F_2) \vee F_2$, we may stipulate that all *regular* worlds other than b are equally plausible, where a world is regular iff it is of the form $\bigsqcup\{\overline{f_m} : m < n\} \sqcup \bigsqcup\{f_m : m \geq n\}$ for some n.

Thus, in revising by F_2, the world with all stones falling is excluded. But it is *not* excluded because it is less *plausible* than the other F_2-worlds. Instead, it does not even come up for consideration at the stage at which the plausiblity filter is applied, because it is not among the worlds that are revisions of b by f_2. Why is it not among those worlds? Because the state f_2 of the second stone falling is taken by Dom to provide no grounds for replacing the state $\overline{f_1}$ of the first stone standing by f_1. Such a reason to change the relevant state obtains only for the other, later stones in the sequence.

So we do not wish to hold that a world w transitions to another v upon revision by s whenever v contains s—this would render our account intensional, and equivalent to AGM. But there are a number of weaker constraints that we may plausibly impose. In particular, for any consistent state s and any world $b \in B$, we shall require that[41]

(1) there is some world w with $s \rightarrow_b w$,
(2) if $s \rightarrow_b w$ then w is a world,
(3) if $s \rightarrow_b w$ then $s \sqsubseteq w$,

[40] A proper definition of such a state, and a proof that it satisfies the assumptions of the case as well as the constraints proposed below, is given in appendix C.

[41] Most of these constraints are similar or identical to ones that Fine imposes on transition relations in his semantics for counterfactuals; cf. (Fine, 2012a: pp. 239ff).

(4) if $s \sqsubseteq b$ then $s \rightarrow_b b$,
(5) if $s \rightarrow_b w$ and $r \sqsubseteq w$, then $s \sqcup r \rightarrow_b w$
(6) if $s \sqcup t \rightarrow_b w$, then $s \rightarrow_b v$ for some $v \leq w$

These constraints, together with the familiar assumptions about plausibility orderings, ensure that filtered revision satisfies natural counterparts of all the basic AGM postulates *except* for *Intensionality*. They also ensure that under the natural interpretation of $\Rightarrow$ in terms of filtered revision, all of Fine's rules from section 4 are valid with the exception of the intensionalist rule of *Substitution*.[42]

The situation is more complicated with respect to the postulates of *Superexpansion* and *Subexpansion*. These are usually stated in a form in which they relate revisions by conjunctions to revisions by their conjuncts. *Superexpansion* then says that $(B * P) \wedge Q$ entails $B*(P \wedge Q)$, and *Subexpansion* says that if $B*P$ is compatible with Q, the converse entailment also holds, so that $B*(P \wedge Q)$ entails $(B*P) \wedge Q$. Now as we have noted before, within an intensional framework, the relation between a conjunction and its conjuncts is simply the relation between a proposition and a proposition entailed by it, and thus the same as the relation between a proposition and a disjunction in which it is a disjunct. As a result, we can also formulate versions of *Superexpansion* and *Subexpansion* that relate revisions by disjunctions to revisions by their disjuncts. These versions will be equivalent to the usual ones under the assumption of intensionally individuated updates, but they will not be equivalent within our hyperintensional framework. We may thus distinguish between the following four principles:

Superexpansion($\wedge$)	$(B * P) \wedge Q$ entails $B * (P \wedge Q)$
Subexpansion($\wedge$)	$B * (P \wedge Q)$ entails $(B * P) \wedge Q$ if $B * P$ is compatible with Q
Superexpansion($\vee$)	$(B * (P \vee Q)) \wedge P$ entails $B * P$
Subexpansion($\vee$)	$B * P$ entails $(B * (P \vee Q)) \wedge P$ if $B * (P \vee Q)$ is compat. with P

It turns out that conditions (1)–(5) on transition relations, together with the conditions on plausibility orderings, ensure that *Superexpansion($\wedge$)* and *Subexpansion($\vee$)* are satisfied. *Superexpansion($\vee$)* and *Subexpansion($\wedge$)* are not in general satisfied.

This is a good thing, though. For as I show in appendix B, there is no way to do so, given the other principles and constraints, without the account collapsing again into AGM and thereby validating *Intensionality*. Moreover, we can construct compelling counter-examples to these postulates on the basis of our example cases. For simplicity,

[42] Proofs, here and below, are again delegated to appendix A.—It is worth mentioning that for the basic AGM postulates, it is sufficient to impose conditions (1)–(4). (5) and (6) are only required for the versions of the supplementary postulates given below, and for Fine's rule of *Transitivity*.

consider the dominos case again. For *Superexpansion($\vee$)*, let P be the proposition that F_2 and Q the proposition that $F_1 \wedge F_2$. Then as we have argued, it is permissible for $B*P$ to rule out that F_1 while $B*(P \vee Q)$—and then also $(B*(P \vee Q)) \wedge P$—does not, and therefore fails to entail $B * P$, in violation of *Superexpansion($\vee$)*. For *Subexpansion($\wedge$)*, let P be as before and let Q be the proposition that $(F_1 \wedge F_2) \vee F_2$. So P says that the second stone fell, and Q says that the second, or the first and the second stone fell. $B*P$ then says the first stone stands, but the second stone and all subsequent ones fell. This is of course compatible with Q; indeed, it entails Q. $P \wedge Q$ is equivalent, even in terms of its truthmakers, to Q. So $B * (P \wedge Q) = B * Q$. But given our assumptions, $B * Q$ makes room for the possibility that all stones fell, and so it cannot entail $B * P$, which does not allow for that possibiliby. A fortiori, $B * Q$ then does not entail $(B * P) \wedge Q$, in violation of *Subexpansion($\wedge$)*.

9. The Advantages of Mightiness

The results of the previous sections show that a viable, hyperintensional theory of rational belief revision can be developed within the framework of truthmaker semantics and on the basis of a conception of the update as mighty. Moreover, we saw that this kind of approach allows us to give a very natural account of what is going on in our puzzle cases, which is much more in line with an intuitive assessment of these cases than any account that could be given within an intensional framework. In this final section of the paper, I want to briefly indicate at a more general and abstract level some of the further advantages of the proposed approach and in particular the use of mighty updates.

The central requirement on a conception of the update is that dynamically inequivalent situations always be assigned distinct updates. As we have seen, there are logically equivalent tm-updates that represent dynamically inequivalent situations. At first glance, if the tm-updates associated with a pair of dynamically inequivalent situations are logically equivalent, it would seem that the pw-updates associated with those situations must be identical. This would show that pw-updates are plainly incapable of capturing the relevant features of occasions for revision. So the question arises how, if at all, intensionalists can avoid this conclusion.

It will be useful to consider a concrete example. Suppose I have hurt my ankle playing football. I take it to be nothing serious but go to the doctor just in case. After examining me, she tells me: ‘Your ankle is sprained, or sprained and broken’. I trust the doctor and see no reason to suspect her to try to mislead me. So I take the situation to tell me that my ankle must be sprained, and that it might in addition be broken. It is then reasonable for me to give up my belief that my ankle is not broken.[43] Now consider a version of the situation in which the doctor tells me simply: 'Your ankle is sprained'. Again, I trust the doctor and see no reason to suspect her to be anything less than fully perspicuous in sharing her opinion of my ankle. So I take the situation to tell me that my ankle is sprained, and I do not take it to tell me that my ankle might be broken. It is then reasonable for me to retain my belief that my ankle is not broken.

Clearly, we have a pair of dynamically inequivalent occasions for revision. Moreover, it seems plausible that under the specified interpretations of the situations, they are to be assigned logically equivalent tm-updates. The tm-update in the first situation—call it the *sprained/broken* scenario—might plausibly be taken to be the truthmaker proposition that my ankle is sprained, or sprained and broken. In the second situation—call it the *sprained* scenario—the tm-update is plausibly taken to be the truthmaker proposition that my ankle is sprained. These propositions, of course, are logically equivalent. Now if the corresponding pw-updates are simply the set of worlds in which these tm-updates are true, then the two situations are assigned the same pw-update, in spite of their dynamic inequivalence. Can the intensionalist plausibly deny the claim that these are the pw-updates?

A natural idea is to point out that the update is supposed to capture the *total* information received by the agent, and that the updates we specified do not satisfy this condition. For example, in the first scenario, I presumably also obtain the information that the doctor assertorically utters the sentence 'Your ankle is sprained, or sprained and broken', and perhaps I obtain the information that the doctor is not convinved that my ankle is not broken. And in the second scenario, I obtain the information that the doctor utters 'Your ankle is sprained' instead, and perhaps take the situation to also tell me that the doctor confidently rules out that my ankle is broken. If we enrich the updates given above by these further bits of information, then the updates assigned to the two situations will not be logically equivalent.[44]

[43] Note that nothing I have said about this scenario depends on it being part of the semantic content of the doctor's utterance that my ankle might be broken. It is perfectly consistent with what I say that this is merely a pragmatic implication. But since I trust the doctor and assume that she is not trying to mislead me, I take on board not only the semantic but also the pragmatic implications of what she says.

[44] While this seems to be the most natural response, it is perhaps not the only possible response. A more comprehensive and detailed examination of the options available here is beyond the scope of this paper, but let me mention one alternative strategy, hinted at in (Spohn, 2014: §6), when discussing a somewhat similar example. The idea is to maintain that in the situation in question, the appropriate

In order to properly evaluate this response, we need to get clearer about the requirement that the update represent the total information received by the agent in the situation under consideration. On the one hand, it is uncontroversial that we need some form of such a completeness requirement: we simply cannot determine the rational responses to a situation purely on the basis of the fact that *part* of what the agent learns is that P, without being told what else the agent learns. On the other hand, a naïvely strict interpretation of the completeness requirement gives rise to severe methodological difficulties. For under such an interpretation, in more or less any realistic situation a doxastic agent might find themselves in, the total information received will be unmanageably rich and complex. For a start, as long as the agent has their eyes open, they would seem to receive, at any point in time, a very rich body of visual information that it is not even feasible to express in words. So if the update needs to capture the total information received in this very demanding sense, we lose the ability to test any proposed theory of belief revision by applying it to realistic scenarios and working out its implications.

In practice, belief revision theorists do not attempt to specify anything like an update that would be complete in this demanding sense. Nor, it might be added, do they normally attempt to fully specify anything like a realistic complete initial belief state that is to be revised, or a complete revised belief state. How is this practice to be justified? To a rough approximation, a natural idea is as follows. First of all, in considering examples, we usually limit attention to the evolution of a certain subset of an agent's beliefs, such as their beliefs concerning the status of certain domino stones or matches, or the whereabouts of their neighbour's twins. We specify those initial beliefs, and tacitly stipulate that in the kind of situation to be considered, any other beliefs the agent might have are irrelevant to how the subset we are considering can rationally be revised. With regard to the update, a related policy is in place: the update is assumed to be complete in the sense of encoding all the information received that is relevant to how the part of the agent's belief system under consideration may rationally be revised. What the example of my injured ankle helps bring out is that the truthmaker approach and the intensional AGM approach differ greatly with respect to how, and how easily, the demands of relevant completeness may be met.

---

response by the agent consists not simply in a revision by some given update, but in a sequence of belief change operations, first simply removing my previous beliefs about the health of my ankle, including the belief that my ankle is not broken, and then revising with the proposition that my ankle is sprained. This suggestion may yield the right results in our example, but absent plausible general principles telling us what situations call for what combinations of operations, the response appears objectionably ad hoc.

Under the truthmaker approach, we can adequately model the example by specifying my initial beliefs about the health of my ankle, and by taking the updates in the two scenarios to be as described above—that my ankle is sprained in the *sprained* scenario, and that my ankle is sprained, or sprained and broken in the *sprained/broken* scenario. Given the assumptions of the example, there seems to be no reason to take these updates to be relevantly incomplete. Under the possible worlds approach, we need to work with a much more complicated model of the situation. In order to capture all the relevant differences about the information received, we have to incorporate in the updates information about which sentences were uttered, or perhaps about which beliefs the doctor holds or does not hold concerning my ankle. To make room for the fact that I can reasonably give up the belief that my ankle is not broken in the *sprained/broken* scenario while retaining the same belief in the *sprained* scenario, we should then say that, roughly speaking, I consider worlds in which the doctor's diagnosis is correct to be more plausible than ones in which it is mistaken.

At least in this kind of case, the truthmaker approach thus affords a simpler, more direct, and more elegant representation of the case. But more significantly, it also allows us to straightforwardly capture intuitive rational constraints that cannot be captured under the alternative, more complicated model. To see this, note that how I am disposed to revise in the *sprained* scenario imposes constraints on how I may rationally be disposed to revise in the *sprained/broken* scenario. In particular, it seems that my beliefs about the health of my ankle should be strictly weaker in the sprained/broken scenario than in the sprained scenario. Under the truthmaker approach, this constraint follows given the logical relationship between the associated tm-updates. For these constitute a pair of the form P and $P \vee (P \wedge Q)$, and we can show using the principles of *Success*, *Consistency* and *Subexpansion($\vee$)* that $B * P$ entails $B * (P \vee (P \wedge Q))$ whenever P is consistent.[45] But it is hard to see how a similar result could be obtained on the basis of a representation of the situations by in terms of the associated pw-updates.[46]

Appendix A. Doxastic States in a Truthmaker Framework

The basic structure of truthmaker semantics is that of a *state-space*, which is a special kind of a partially ordered set. Recall that a partial order on a set S is a two-place

[45] By *Success*, $B * (P \vee (P \wedge Q))$ entails $P \vee (P \wedge Q)$, and hence P. If P is consistent, so is $P \vee (P \wedge Q)$ and thus by *Consistency*, $B * (P \vee (P \wedge Q))$, so $B * (P \vee (P \wedge Q))$ is compatible with P. By *Subexpansion($\vee$)*, it then follows that $B * P$ entails $B * (P \vee (P \wedge Q))$.

[46] [Acknowledgements removed for blind review.]

relation $\sqsubseteq$ that is reflexive—$s \sqsubseteq s$ for all $s \in S$—, transitive—$s \sqsubseteq t$ and $t \sqsubseteq u$ implies $s \sqsubseteq u$ for all $s, t, u \in S$—, and anti-symmetric—$s \sqsubseteq t$ and $t \sqsubseteq s$ implies $s = t$ for all $s, t \in S$. Call a partial order $\sqsubseteq$ on a set S *complete* iff for every subset T of S, a least upper bound exists, i.e. there is an element $s \in S$ such that $t \sqsubseteq s$ for all $t \in T$, and $s \sqsubseteq u$ whenever $t \sqsubseteq u$ for all $t \in S$. By designating a certain subset of the states as the set of *possible* or *consistent* states, we obtain a modalized state-space.

Definition 1. *A* modalized state-space *is a triple* $(S, S^\Diamond, \sqsubseteq)$ *such that*

(1) S *is a non-empty subset,*
(2) $\sqsubseteq$ *is a complete partial order on* S*, and*
(3) $S^\Diamond$ *is a non-empty subset of* S *such that* $s \in S^\Diamond$ *whenever* $s \sqsubseteq t$ *and* $t \in S^\Diamond$.

Informally, the members of S are the states, $\sqsubseteq$ is the parthood-relation, and $S^\Diamond$ is the set of possible, or consistent states. For $T \subseteq S$, we write $\bigsqcup T$ for the least upper bound of T, which we also call the *fusion* of the members of T, and often write as $t_1 \sqcup t_2 \sqcup \ldots$ when $T = \{t_1, t_2, \ldots\}$. We call a *(possible) world* any possible state that contains every state it is compatible with, and we denote their set by S^w. A modalized state-space is called a *W-space* iff every possible state is part of a possible world, and it is called *topsy* if it contains only one impossible state. This will then be the one state, written $\blacksquare$, which contains every state. Throughout this section and the next, we shall be working within some fixed, topsy W-space $\mathcal{S} = (S, S^\Diamond, \sqsubseteq)$.

A *proposition* P is any non-empty subset of S. A proposition is consistent iff one of its members is, and propositions P and Q are compatible iff some member of P is compatible with some member of Q. The conjunction $P \wedge Q$ of unilateral proposition P and Q is $\{p \sqcup q : p \in P \text{ and } q \in Q\}$. Note that this is non-empty even if P and Q are incompatible, in which case $P \wedge Q$ is $\{\blacksquare\}$.[47] The disjunction $P \vee Q$ is $P \cup Q$. A proposition P is said to *(loosely[48]) entail* ($\models$) a proposition Q iff every world containing a truthmaker of P contains a truthmaker of Q .

We now turn to the task of defining the class of permissible doxastic states. We do this using a notion of a *coherent pair* of a *plausibility* ordering and a *transition* relation.

Definition 2. *A* plausibility ordering *is a two-place relation* $\leq$ *on* $S^w \cup \{\blacksquare\}$ *satisfying the following conditions, where* $g(X) := \{w \in X : w \leq v \text{ for all } v \in X\}$ *for all* $X \subseteq S^w \cup \{\blacksquare\}$*:*

[47] The main reason for countenancing the impossible $\blacksquare$, in the present context, is to avoid technical inconveniences such as having to allow for empty propositions.

[48] There are several other, narrower relations of entailment that may be defined within the truthmaker framework. For present purposes, however, we may confine attention to this one.

(P-Connectedness)	$w \leq v$ *or* $v \leq w$
(P-Transitivity)	$w \leq v$ *and* $v \leq u$ *implies* $w \leq u$
(P-Limit)	$g(X) \neq \emptyset$ *if* $\emptyset \subset X \subseteq S^w$
(P-Inconsistency)	$\blacksquare \leq s$ *implies* $s = \blacksquare$

These are exactly the conditions imposed under the possible worlds approach, except for the added clause dealing with $\blacksquare$. Given any plausibility ordering, we often use B to refer to $g(S^w)$, since this is the set of worlds at which the agent's beliefs are true.

Definition 3. *A* transition relation *is a three-place relation $\rightarrow$ on S subject to the following conditions:*

(T-Success)	$s \rightarrow_u t$ *implies* $t \sqsupseteq s$
(T-Completeness)	*if* $u \in S^w$ *and* $s \rightarrow_u t$ *then* $t \in S^w \cup \{\blacksquare\}$
(T-Consistency)	*if* $s, u \in S^\diamond$ *then* $s \rightarrow_u t$ *for some* $t \in S^\diamond$
(T-Vacuity)	*if* $s \sqsubseteq u$ *then* $s \rightarrow_u u$
(T-Incorporation)	*if* $s \rightarrow_u t$ *and* $r \sqsubseteq t$ *then* $s \sqcup r \rightarrow_u t$

Definition 4. *Let $\leq$ and $\rightarrow$ be a pair of plausibility ordering and transition relation. The operations of* wayward revision $\circ$ *and* (filtered) revision $*$ *induced by $\leq$ and $\rightarrow$ are defined as follows, for P a non-empty subset of S:*

$$B \circ P := \{t \in S : p \rightarrow_b t \text{ for some } p \in P \text{ and } b \in B\}$$
$$B * P := g(B \circ P)$$

Definition 5. *A pair of a plausibility ordering $\leq$ and a transition relation $\rightarrow$ is* coherent *iff, whenever $b \in B$:*

(PT-Existence)	$s \rightarrow_b t$ *for some* $t \in S$
(PT-Link)	*if* $w \in B \circ \{s \sqcup t\}$ *then* $v \leq w$ *for some* $v \in B \circ \{s\}$

Theorem 1. *Let $*$ be the revision function induced by some coherent pair of plausibility ordering and transition relation. Then*

(R-Success)	$B * P \models P$
(R-Vacuity)	$B * P \models B$ *if* B *is compatible with* P
(R-Inclusion)	$B \wedge P \models B * P$
(R-Consistency)	$B * P$ *is consistent if* P *is*
(R-Superexpansion($\wedge$))	$(B * P) \wedge Q \models B * (P \wedge Q)$
(R-Subexpansion($\vee$))	$B * P \models (B * (P \vee Q)) \wedge P$ *if* $B * (P \vee Q)$ *is compatible with* P

Proof. Note that $B * P \subseteq S^w \cup \{\blacksquare\}$ for all non-empty P. So to establish that any revision entails some proposition Q, we need to show that every truthmaker of the revision contains a truthmaker of Q as part.

(R-Success): If $s \in B * P$, then $p \to_b s$ for some $p \in P$ and $b \in B$, and by (T-Success), $s \sqsupseteq p$.

(R-Vacuity): Suppose B is compatible with P, and let $b \in B$ be compatible with $p \in P$. Since b is a world, it follows that $b \sqsupseteq p$, so by (T-Vacuity) $p \to_b b$, and hence $b \in B \circ P$. Since $b \in g(S^w)$, it follows that $b \leq v$ for all $v \in B \circ P$ and hence $b \in B * P$. But then $v \leq b$ for all $v \in B * P$, so $B * P \subseteq B$ and hence $B * P \models B$.

(R-Inclusion): Note that by (PT-Existence), $B * P$ is non-empty. Suppose $s \in B \wedge P$, and let $b \in B$ and $p \in P$ be such that $s = b \sqcup p$. If $s = \blacksquare$, s contains every state, and hence some verifier of $B * P$. If s is consistent, then b is compatible with p, so since $b \in S^w$, $b \sqsupseteq p$, and thus $s = b$. By (T-Vacuity), $p \to_b b$, $b \in B \circ P$. Since $b \in B = g(S^w)$, $b \leq w$ for all $w \in B \circ P$, hence $s = b \in B * P$.

(R-Consistency): Suppose $s \in P$ is consistent and let $b \in B$. By (T-Consistency), $v \in S^w$ for some v with $s \to_b v$, so $B \circ P$ has a consistent member. By (P-Inconsistency), it follows that $g(B \circ P)$ has some (indeed, only) consistent members.

(R-Superexpansion($\wedge$)): Note first that by (PT-Existence), $B * (P \wedge Q)$ is non-empty. Now suppose $s \in (B * P) \wedge Q$. Then $s = t \sqcup q$ for some $t \in B * P$ and $q \in Q$. If $s = \blacksquare$, s contains every state, and so some verifier of $B * (P \wedge Q)$. If s is consistent, then t is a possible world, and hence contains q as part. Since $t \in B * P$, $p \to_b t$ for some $p \in P$ and $b \in B$. By (T-Incorporation), also $p \sqcup q \to_b t$, and hence $t \in B \circ (P \wedge Q)$. Now consider any $v \in B \circ (P \wedge Q)$. Let $p' \in P$ and $q' \in Q$ be such that $v \in B \circ \{p' \sqcup q'\}$. Then by (PT-Link) there is some $u \leq v$ with $u \in B \circ \{p\}$ and hence $u \in B \circ P$. Since $t \in B * P$, it follows that $t \leq u$ and hence by (P-Transitivity) $t \leq v$. So $t \in g(B \circ (P \wedge Q)) = B * (P \wedge Q)$, as desired.

(R-Subexpansion($\vee$)): Suppose $B * (P \vee Q)$ is compatible with P. Suppose $s \in B * P$. As before, the case in which s is inconsistent is easy. Suppose instead $s = w \in S^w$ and $p \to_b w$ with $p \in P$, and $b \in B$. We wish to show that $w \in (B * (P \vee Q)) \wedge P$. Since $w \sqsupseteq p$ and $p \in P$, it suffices to show that $w \in B * (P \vee Q)$. Since $w \in B * P$, $w \in B \circ P$ and hence $w \in B \circ (P \vee Q)$. It remains to show that $w \leq v$ for all $v \in B \circ (P \vee Q)$. Now note that since $B * (P \vee Q)$ is compatible with P, there is some $w' \in g(B \circ (P \vee Q))$ with $w' \leq v$ for all $v \in B \circ (P \vee Q)$, and $w' \sqsupseteq p'$ for some $p' \in P$. It then suffices to show that $w \leq w'$. Now either $w' \in B \circ P$ or $w' \in B \circ Q$. If $w' \in B \circ P$, then $w \leq w'$ is immediate from $w \in g(B \circ P)$. So suppose $w' \in B \circ Q$, so $q \to_b w'$ for some $q \in Q$. Since $w' \sqsupseteq p'$, by (T-Incorporation), $p' \sqcup q \rightarrow_b w'$. Then by (PT-Link), $u \in B \circ \{p'\}$, so $u \in B \circ P$, and hence $w \leq u$. By (P-Transitivity), $w \leq w'$, as desired. □

(R-Success), (R-Vacuity), (R-Inclusion), and (R-Consistency) are the obvious counterparts in our (semantic) setting to the (syntactically formulated) AGM postulates of *Success*, *Vacuity*, *Inclusion* and *Consistency*. The postulate of *Closure* serves mainly to ensure intensionality with respect to belief states, which is guaranteed under our account by the identification of belief states with the set of possible worlds at which they are true. The *Intensionality* postulate, of course, does not hold. Within our semantic setting, the only valid version of this principle is the triviality that $B * P = B * Q$ if $P = Q$. Under a syntactic formulation of the theory, though, we would have the non-trivial principle that $K * \alpha = K * \beta$ if α and β are *exactly equivalent*, i.e. have the same exact truthmakers.[49]

Moreover, as expected, all Finean rules from section 4 except for the intensionalist rule of *Substitution* are valid under the obvious interpretation of $\Rightarrow$.

Theorem 2. *Let $*$ be the revision function induced by some coherent pair of plausibility ordering and transition relation. For any propositions P, Q, let $P \Rightarrow Q$ hold iff $B * P \models Q$. Then*

(R-Entailment)	$P \Rightarrow Q$ *whenever* $P \models Q$
(R-Transitivity)	*If* $P \Rightarrow Q$ *and* $P \wedge Q \Rightarrow R$ *then* $P \Rightarrow R$
(R-Conjunction)	*If* $P \Rightarrow Q$ *and* $P \Rightarrow R$ *then* $P \Rightarrow Q \wedge R$
(R-Disjunction)	*If* $P \Rightarrow R$ *and* $Q \Rightarrow R$ *then* $P \vee Q \Rightarrow R$

Proof. (R-Entailment) and (R-Conjunction) are immediate from the definition of $\Rightarrow$ and the fact that the $\models$-consequences of a proposition are closed under conjunction. (R-Disjunction) is immediate from the observation that $g(X \cup Y) \subseteq g(X) \cup g(Y)$.

(R-Transitivity): Assume $B * P \models Q$ and $B * (P \wedge Q) \models R$. If P is inconsistent, $P \Rightarrow R$ follows immediately given (Success). So suppose P is consistent. Then $B * P \subseteq S^w$. So let $w \in B * P$, and suppose $p \rightarrow_b w$ with $p \in P$ and $b \in B$. We need to show that $w \sqsupseteq r$ for some $r \in R$. Since $B * P \models Q$, we have $w \sqsupseteq q$ for some $q \in Q$. Then $p \sqcup q$ is consistent. By (T-Incorporation), $p \sqcup q \rightarrow_b w$, so $w \in B \circ (P \wedge Q)$. Now let $v \in B \circ (P \wedge Q)$, and let $p' \in P$ and $q' \in Q$ be such that $v \in B \circ \{p' \sqcup q'\}$. By (PT-Link), $u \in B \circ \{p'\}$ and hence $u \in B \circ P$ for some $u \leq v$. But since $w \in B * P$, $w \leq u$ and hence $w \leq v$. So $w \in B * (P \wedge Q)$. Since $B * (P \wedge Q) \models R$, $w \sqsupseteq r$ for some $r \in R$, as desired. □

---

49 On the logic of this equivalence relation, see Fine (2016); Correia (2016); Krämer (202x).

Appendix B. Expansion and Collapse

We now show that the validity of *Subexpansion(∧)* or *Superexpansion(∨)* would collapse our account into intensional AGM. Some additional notation will be helpful. For any proposition Q, let Q^w be the set of Q-worlds, i.e. $\{w \in S^w : w \sqsupseteq q$ for some $q \in Q\}$. The AGM revision of B by P is then simply $g(P^w)$. Now consider the following constraint on the connection between plausibility orderings and transition relations:

(PT-Expansion) $\quad g(\{p\}^w) \subseteq B \circ \{p\}$

It turns out that whenever (PT-Expansion) is satisfied, the resulting revision function is an AGM revision function.[50]

Proposition 3. *Let $\leq$ and $\rightarrow$ be a coherent pair of plausibility ordering and transition relation that satisfy (PT-Expansion). Then $B * P = g(P^w)$ whenever P is consistent.*

Proof. Let P be consistent. Then $B * P = B * (P \cap S^\diamond)$ and $g(P^w) = g((P \cap S^\diamond)^w)$, so we may assume without loss of generality that $P \subseteq S^\diamond$.

Suppose first that $w \in g(P^w)$. Let $p \in P$ be such that $w \sqsupseteq p$. Then $w \in g(\{p\}^w)$, so by (PT-Expansion), $w \in B \circ \{p\}$ and hence $w \in B \circ P$. Suppose $v \in B \circ P$. Then $v \in P^w$, so since $w \in g(P^w)$, $w \leq v$. It follows that $w \in B * P$.

Suppose now that $w \in B * P$, so $w \in g(B \circ P)$. Then $w \in P^w$. Now let $v \in P^w$. Pick a world $u \in g(P^w)$, so $u \leq v$, and let $p' \in P$ be such that $u \sqsupseteq p'$ and hence $u \in g(\{p\}^w)$. By (PT-Expansion), $u \in B \circ \{p\}$ and hence $u \in B \circ P$. Since $w \in g(B \circ P)$, it follows that $w \leq u$ and hence $w \leq v$. So $w \in g(P^w)$, as desired. □

Any violation of (PT-Expansion), however, yields a violation of both *Subexpansion(∧)* and *Superexpansion(∨)*.

Proposition 4. *Let $\leq$ and $\rightarrow$ be a coherent pair of plausibility ordering and transition relation. Let $w \in S^w$, $p \sqsubseteq w$, and $w \leq v$ for all $v \in B \circ \{p\}$, but $w \notin B \circ \{p\}$. Then both* Subexpansion(∧) *and* Superexpansion(∨) *are invalid.*

Proof. For *Subexpansion(∧)*, set $Q = \{p, w\}$ and $P = \{p\}$. It suffices to show that (i) $w \notin B * P$, (ii) $B * P$ is compatible with Q, and (iii) $w \in B * (P \wedge Q)$. But (i) is immediate from the assumption that $w \notin B \circ \{p\}$. For (ii), note that since p is consistent, by (T-Consistency) and (T-Completeness), for some world u, $u \in B * P$, and by (T-Success), $u \sqsupseteq p$, so $B * P$ is compatible with Q. For (iii), note that $B \circ (P \wedge Q) = B \circ \{p, w\} =$

50 Modulo some irrelevant differences resulting from the subtly different treatment of inconsistent updates using ■.

$(B \circ \{p\}) \cup (B \circ \{w\})$. Since w is a world, by (T-Success) and (T-Consistency), $w \rightarrow_b u$ implies $u = w$ for all $b \in B$, so $B \circ \{w\} = \{w\}$. It follows that $w \in B \circ (P \wedge Q)$. Moreover, for all $v \in B \circ (P \wedge Q)$, either $v = w$ or $v \in B \circ \{p\}$. By assumption, either way we have $w \leq v$, and hence $w \leq v$ for all $v \in B \circ (P \wedge Q)$, so $w \in B * (P \wedge Q)$.

For *Superexpansion(∨)*, set $Q = \{w\}$ and $P = \{p\}$. As before, $w \notin B * P$. So it suffices to show that $w \in B * (P \vee Q)$, since then by $w \sqsupseteq p$ also $w \in (B * (P \vee Q)) \wedge P$. But $P \vee Q = \{p, w\}$, and we already showed above that $w \in B * \{w, p\}$. □

It is natural to wonder if the collapse may be avoided by invalidating *Subexpansion(∨)* and *Superexpansion(∧)* instead of *Subexpansion(∧)* and *Superexpansion(∨)*. But this is not so. Indeed, (PT-Expansion) implies (PT-Link), the principle required to establish *Subexpansion(∨)* or *Superexpansion(∧)*. For assume $s \sqcup t \rightarrow_b w$. If $s \sqcup t$ is inconsistent, $w = \blacksquare$ and $w \leq v$ for any $v \in S^w$. If $s \sqcup t$ is consistent, w is a world containing $s \sqcup t$ and hence $w \in \{s\}^w$. By (PT-Expansion) $g(\{s\}^w) \subseteq B \circ \{s\}$. By (P-Limit), $g(\{s\}^w)$ is non-empty, so let $v \in g(\{s\}^w)$. Then $v \leq w$ and $v \in B \circ \{s\}$ as required.

Appendix C. A Space for Dominos

Finally, we construct a model of a doxastic state that satisfies the assumptions of the domino case and whose revision function is obtained from a coherent pair of plausibility ordering and transition relation. Let f_1, f_2, ... be a countable infinity of sentence letters, and let L be the corresponding set of literals, i.e. the set including exactly the sentence letters f_n as well as their negations, which we write as $\overline{f_n}$. Say that a subset s of L is consistent iff for all n, at most one of f_n and $\overline{f_n}$ is a member of s. Let $S^\diamond = \{s \subseteq L : s$ is consistent$\}$ and $S = S^\diamond \cup \{L\}$. It is straightforward to show that $\mathcal{S} = (S, S^\diamond, \sqsubseteq)$, with $\sqsubseteq$ interpreted as the subset-relation, is a topsy W-space, the set of worlds S^w being the set of the maximal consistent subsets of L.

We now define a plausibility ordering $\leq$ on the worlds. First, let $b = \{\overline{f_n} : n \in N\}$. Say that a world w is *regular* if for some n, $w = \{\overline{f_m} : m < n\} \cup \{f_m : m \geq n\}$, and irregular if not regular and distinct from b. Then for w, $v \in S^w$, we let $w \leq v$ iff (a) $w = b$, or (b) w is regular and $v \neq b$, or (c) w is irregular and v is irregular or identical to $\blacksquare$, or (d) $w = v = \blacksquare$. It is readily verified that $\leq$ satisfies the conditions of (P-Connectedness), (P-Transitivity), (P-Limit), and (P-Inconsistency).

Next, we define a transition relation $\rightarrow$. It will be sufficient to specify revisions of b by any state s. Indeed, for each state s we shall always specify a *unique* revision of b by s. If s is $\blacksquare$, we let $s \rightarrow_b t$ iff $t = \blacksquare$. If s is consistent, then for each n, we consider the largest $m \leq n$, if any, for which either f_m or $\overline{f_m}$ is a member of s, and we include f_n in our output state if s contains f_m, and $\overline{f_n}$ otherwise. More precisely, say that s is n-positive iff (a) $f_m \in s$ for some $m \leq n$, and (b) $f_m \in s$ for m the greatest number $\leq n$ for which either $f_m \in s$ or $\overline{f_m} \in s$. If s is not n-positive, then it is n-negative. Then let $\phi(s,n) = f_n$ if s is n-positive and $\overline{f_n}$ otherwise, and for consistent s, let $s \rightarrow_b t$ iff $t = \{\phi(s,n) : n \in N\}$.

Proposition 5. *$\leq$ and $\rightarrow$ are a coherent pair of plausibility ordering and transition relation on $\mathcal{S}$.*

Proof. We skip the straightforward proof that $\leq$ is a plausibility ordering.

(T-Success): if $s = \blacksquare$, then $s \rightarrow_b t$ implies $t = \blacksquare$ and hence $t \sqsupseteq s$. If s is consistent, then $s \rightarrow_b t$ implies $t = \{\phi(s,n) : n \in N\}$. Suppose $f_n \in s$. Then s is n-positive and hence $f_n \in t$. Suppose $\overline{f_n} \in s$. Since s is consistent, $f_n \notin s$. So s is not n-positive, and hence $\overline{f_n} \in t$. So $s \sqsubseteq t$, as required.

(T-Completeness): $s \rightarrow_b t$ implies that either $t = \blacksquare$ or $t = \{\phi(s,n) : n \in N\}$. By construction, for each n, $\{\phi(s,n) : n \in N\}$ has either f_n or $\overline{f_n}$ as a member, so $\{\phi(s,n) : n \in N\} \in S^w$.

(T-Consistency): $s \rightarrow_b t$ implies $t = \{\phi(s,n) : n \in N\}$ given that s is consistent. By construction of $\{\phi(s,n) : n \in N\}$, it follows that for all n, at most one of f_n and $\overline{f_n}$ are members of $\{\phi(s,n) : n \in N\}$, so t is consistent.

(T-Vacuity): If $s \sqsubseteq b$, then s contains no f_n as member. By construction, then neither does $\{\phi(s,n) : n \in N\}$, so $\{\phi(s,n) : n \in N\} = b$.

(T-Incorporation): Assume $s \rightarrow_b t$ and $r \sqsubseteq t$. We need to show that $s \sqcup r \rightarrow_b t$. Suppose first that s is inconsistent. Then $s \sqcup r = s$, and $s \sqcup r \rightarrow_b t$ follows immediately. Suppose then that s is consistent, so $t = \{\phi(s,n) : n \in N\}$. By definition of $\rightarrow$, $s \sqcup r \rightarrow_b \{\phi(s \sqcup r, n) : n \in N\}$, so it suffices to show that for all n, s is n-positive iff $s \sqcup r$ is.

Suppose first that $s \sqcup r$ is n-positive. So for m the greatest number $\leq n$ such that $s \sqcup r$ contains either f_m or $\overline{f_m}$, we have $f_m \in s \sqcup r$. Since $s \sqcup r \sqsubseteq t = \{\phi(s,n) : n \in N\}$, it follows that $\phi(s,m)$ is f_m, so s is m-positive. But since m is the greatest number $\leq n$ for which $s \sqcup r$ contains either f_m or $\overline{f_m}$, it follows that there can be no number k between m and n for which s contains $\overline{f_k}$, and hence it follows that s is n-positive also.

Suppose now that $s \sqcup r$ is n-negative. Then either (a) there is no $m \leq n$ with $f_m \in s \sqcup r$, or (b) we have $\overline{f_m} \in s \sqcup r$ for m the greatest number $\leq n$ for which either $f_m \in s \sqcup r$ or $\overline{f_m} \in s \sqcup r$. If (a), then there is no $m \leq n$ with $f_m \in s$, so s is n-negative. If (b), then since $s \sqcup r \sqsubseteq t = \{\phi(s,n) : n \in N\}$, it follows that $\phi(s,m)$ is $\overline{f_m}$, so s is m-negative. But since m is the greatest number $\leq n$ for which $s \sqcup r$ contains either f_m or $\overline{f_m}$, it follows that there can be no number k between m and n for which s contains f_k, and hence it follows that s is n-negative also.

(PT-Existence): note that $g(S^w) = \{b\}$, and $s \rightarrow_b \{\phi(s, n) : n \in N\}$ for all $s \in S^\Diamond$, and $s \rightarrow_b \blacksquare$ for $s = \blacksquare$.

(PT-Link): Suppose $v \in B \circ \{s \sqcup r\}$, so $s \sqcup t \rightarrow_b v$. We need to show that $s \rightarrow_b u$ for some $u \leq v$. To that end, we establish

(1) if $s \sqcup t \rightarrow_b b$, then $s \rightarrow_b b$
(2) if $s \sqcup t \rightarrow_b w$ for some regular w, then $s \rightarrow_b v$ for some regular v
(3) If $s \sqcup t \rightarrow_b w$ for some irregular, consistent w, then $s \rightarrow_b v$ for some consistent v

For (1), note that if $s \sqcup t \rightarrow_b b$, then $s \sqcup t \sqsubseteq b$ by (T-Success), so $s \sqsubseteq b$, so $s \rightarrow_b b$ by (T-Vacuity). For (2), we prove the contrapositive. Suppose that $s \rightarrow_b v$ and $v = \{\phi(s, n) : n \in N\}$ is irregular. By definition of irregularity there are $m < k$ with both f_m and $\overline{f_k}$ members of v. By definition of ϕ and the fact that $b = \{\overline{f_n} : n \in N\}$ it follows that for some $m' < k'$, both $f_{m'}$ and $\overline{f_{k'}}$ are members of s, and hence of $s \sqcup t$. By (T-Success), both $f_{m'}$ and $\overline{f_{k'}}$ are members of w, which rules out w being regular. For (3), note that if $s \sqcup t \rightarrow_b w$ with w consistent, then $s \sqcup t$ is consistent, hence so is s, and so by consistency so is v with $s \rightarrow_b v$. □

We now show that our doxastic state satisfies the assumptions of the domino case under their obvious interpretation. Since some of these assumptions concern negated propositions, we will move to a bilateral conception of propositions as a pair of set of truthmakers and a set of falsitymakers. We shall take the revision of a belief state by a bilateral proposition to be simply the revision by the set of truthmakers, so our overall account of revision is not changed.

More precisely, we call a *bilateral proposition* $\mathbf{P}$ any pair of unilateral propositions. The first (second) coordinate of $\mathbf{P}$ is denoted by $\mathbf{P}^+$ ($\mathbf{P}^-$) and comprises the truthmakers (falsitymakers) of $\mathbf{P}$. Let $\mathbf{P} \wedge \mathbf{Q} = (\mathbf{P}^+ \wedge \mathbf{Q}^+, \mathbf{P}^- \vee \mathbf{Q}^-)$, $\mathbf{P} \vee \mathbf{Q} = (\mathbf{P}^+ \vee \mathbf{Q}^+, \mathbf{P}^- \wedge \mathbf{Q}^-)$, and $\neg\mathbf{P} = (\mathbf{P}^-, \mathbf{P}^+)$. $\mathbf{P}$ is said to be *exhaustive* iff every $w \in S^w$ contains either a member of $\mathbf{P}^+$ or a member of $\mathbf{P}^-$ as a part, and it is said to be *exclusive* iff no $w \in S^w$ contains both a member of $\mathbf{P}^+$ and a member of $\mathbf{P}^-$ as a part. Both properties can be shown to be preserved under the boolean operations, and it can also be shown that the logic of loose entailment over exclusive and exhaustive propositions is classical (cf. (Fine, 2017a: pp. 665ff)).

Now let $\mathbf{F}_n = (\{f_n\}, \{\overline{f_n}\})$ and $B = g(S^w)$. Note that $\mathbf{F}_n$ is always exclusive and exhaustive. Let $\mathbf{P} \Rightarrow \mathbf{Q}$ hold iff $B * \mathbf{P}^+ \models \mathbf{Q}^+$, and let $\Rightarrow \mathbf{Q}$ hold iff $B \models \mathbf{Q}^+$. Then

Proposition 6. *The belief state and revision function induced by $\leq$ and $\rightarrow$ satisfy the assumptions of the domino case:*

(B) $\Rightarrow \neg \mathbf{F}_n$ *for all n*

(D.+) $\mathbf{F}_n \Rightarrow \mathbf{F}_m$ *for any* $m \geq n$

(D.−) $\mathbf{F}_n \Rightarrow \neg \mathbf{F}_m$ *for any* $m < n$

Proof. (B) is immediate from the facts that $B = \{b\}$ and the definition $b = \{\overline{f_n} : n \in N\}$.

For (D.+) and (D.−), note that f_n is m-positive iff $m \geq n$, so $f_n \rightarrow_b t$ iff $t = \{\overline{f_m} : m < n\} \cup \{f_m : m \geq n\}$. So $B * \mathbf{F}_n^+ = \{\{\overline{f_m} : m < n\} \cup \{f_m : m \geq n\}\}$. So the sole truthmaker of $B * \mathbf{F}_n^+$ contains a truthmaker of $\mathbf{F}_m$ whenever $m \geq n$, as required for (D.+), and a truthmaker of $\neg \mathbf{F}_m$ whenever $m < N$, as required for (D.−). □

References

Alchourrón, C. E., P. Gärdenfors, and D. Makinson (1985). On the Logic of Theory Change: Partial Meet Contraction and Revision Functions. *Journal of Symbolic Logic 50*, 510–530.

Berto, F. (2019). Simple Hyperintensional Belief Revision. *Erkenntnis 84*, 559–575.

Correia, F. (2016). On the Logic of Factual Equivalence. *The Review of Symbolic Logic 9(1)*, 103–122.

Embry, B. (2014). Counterfactuals Without Possible Worlds? A Difficulty for Fine's Exact Semantics for Counterfactuals. *The Journal of Philosophy 111*, 276–287.

Fagin, R. and J. Y. Halpern (1988). Belief, Awareness, and Limited Reasoning. *Artificial Intelligence*, 39–76.

Fermé, E. and S. O. Hansson (2018). *Belief Change. Introduction and Overview*. Springer.

Fine, K. (2012a). Counterfactuals without Possible Worlds. *The Journal of Philosophy 109(3)*, 221–246.

Fine, K. (2012b). A Difficulty for the Possible Worlds Analysis of Counterfactuals. *Synthese 189*, 29–57.

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

Fine, K. (2017a). A Theory of Truthmaker Content I: Conjunction, Disjunction and Negation. *Journal of Philosophical Logic 46(6)*, 625–674.

Fine, K. (2017b). A Theory of Truthmaker 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. (ms). Review of Steve Yablo's Aboutness. Unpublished manuscript.

Gärdenfors, P. (1986). Belief Revision and the Ramsey Test for Conditionals. *Philosophical Review 95*, 81–93.

Gärdenfors, P. (1987). Variations on the Ramsey Test: More Triviality Results. *Studia Logica 46*, 319–325.

Gärdenfors, P. (1988). *Knowledge in Flux. Modeling the Dynamics of Epistemic States*. MIT Press.

Grove, A. (1988). Two Modellings for Theory Change. *Journal of Philosophical Logic 17*, 157–170.

Hansson, S. O. (1999). *A Textbook of Belief Dynamics. Theory Change and Database Updating*. Kluwer.

Hansson, S. O. (2017). Logic of Belief Revision. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2017 ed.). Metaphysics Research Lab, Stanford University.

Huber, F. (2013). Belief Revision I: The AGM Theory. *Philosophy Compass 8/7*, 604–612.

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

Katsuno, H. and A. O. Mendelzon (1991). Propositional knowledge base revision and minimal change. *Artifical Intelligence 52*, 263–294.

Kourousias, G. and D. Makinson (2007). Parallel Interpolation, Splitting, and Relevance in Belief Change. *Journal of Symbolic Logic 72(3)*, 994–1002.

Krämer, S. (202x). Truthmaker Equivalence. To appear in a volume for Alan Weir.

Lewis, D. K. (1973). *Counterfactuals*. Blackwell.

Makinson, D. (1990). The Gärdenfors Impossibility Theorem in Nonmonotonic Contexts. *Studia Logica 49*, 1–6.

Makinson, D. (2009). Propositional Relevance through Letter-Sharing. *Journal of Applied Logic 7(4)*, 377–387.

Parikh, R. (1999). Beliefs, Belief Revision, and Splitting Languages. *Logic, Language, and Computation 96(2)*, 266–278.

Peppas, P., S. Chopra, and N. Foo (2004). Distance semantics for relevance-sensitive belief revision. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Ninth International Conference (KR2004)*, pp. 319–328.

Rantala, V. (1982). Impossible World Semantics and Logical Omniscience. *Acta Philosophica Fennica*, 18–24.

Rothschild, D. and S. Yablo (202x). Permissive Updates. To appear in a volume for Kit Fine.

Spohn, W. (2014). AGM, Ranking Theory, and the Many Ways to Cope with Examples. In S. O. Hansson (Ed.), *David Makinson on Classical Methods for Non-Classical Problems*, pp. 95–118. Springer.

Stalnaker, R. (1968). A Theory of Conditionals. In N. Rescher (Ed.), *Studies in Logical Theory*, Volume 2, pp. 98–112. Blackwell.

Stalnaker, R. (2009). Iterated Belief Revision. *Erkenntnis 70(2)*, 189–209.

van Fraassen, B. (1969). Facts and Tautological Entailment. *Journal of Philosophical Logic 66(15)*, 477–487.

von Wright, G. H. (1963). *Norm and Action*. Routledge & Kegan.

Yablo, S. (2011). A Problem about Permission and Possibility. In A. Egan and B. Weatherstone (Eds.), *Epistemic Modality*, pp. 270–294. Oxford University Press.

Yablo, S. (2015). *Aboutness*. Princeton University Press.

Özgün, A. and F. Berto (2020). Dynamic Hyperintensional Belief Revision. *The Review of Symbolic Logic*, 1–46.