Modal Communication

Sam Carter

Abstract

This paper develops a novel account of communication with epistemic modals. §2 introduces a range of conditions on modal communication. §§3-4 provide a taxonomy of existing theories and identify important limitations on their adequacy. §5 introduces a new kind of theory, which supplements a broadly informational approach to epistemic modality with a generalization of AGM revision. In addition to demonstrating that this theory satisfies each of the conditions on modal communication identified in §2, I prove a representation theorem for the revision operation and show that it validates a generalized version of the Levi Identity.

1 Introduction

Communication, at an appropriate level of abstraction, can be understood as information change. For a core vocabulary of basic expressions, we have a simple and seemingly satisfying theory of how information is exchanged in conversation. An agent utters a sentence which she accepts and which encodes the message she wishes to communicate in its content. Her audience—assuming they are credulous—will then modify their own information on the basis of this content, with the result that they come to accept the sentence uttered.

Elegant as it is, this theory runs into difficulty explaining the communicative effects of epistemic modals. Even when overtly focused on facts about communication, existing theories of epistemic modality have largely failed to address the way agents update in response to learning about what must or may be the case. (For examples of sophisticated theories which do address this issue, see in particular Hawke & Steinert-Threlkeld (2018) and Rothschild & Yablo (2021), discussed in §3.2.2).

This paper sets out to extend the simple theory of communication to a modal language. It has three parts. The first part identifies a number of descriptive and theoretical conditions which it would be nice for a theory of modal communication to satisfy (§2). The second, negative, part of the paper presents an informal taxonomy of existing theories (§3), and provides proofs in a more formal setting that neither dynamic theories nor a common sub-class of static theories can satisfy all of the desiderata (§4). The third part of the paper develops the positive proposal. I show (§5) that it is possible to develop a form of static theory which satisfies each of the desiderata in §2 while avoiding the issues in §4. This theory combines a conservative generalization of AGM revision (Alchourrón *et al.* (1985)) with a family of theories of epistemic modals which agree on some basic constraints (Veltman (1996); Yalcin (2007, 2011); Santorio (2018); Hawke & Steinert-Threlkeld (2018, 2020); Gillies (2020); Goldstein (2021); Rothschild & Yablo (2021)). §6 is the conclusion.

2 Epistemic Modals in Communication

There's been a theft at the seminary. Based on your initial investigation, you think that it was either the Abbot or the Bishop who did it. However, before accusing anyone, you enlist the services of a professional detective. Consider a (non-exhaustive) list of things the detective could tell you:

(1) It was the Abbot.

(2) It must have been the Abbot.

(3) It can't have been the Bishop.

(4) It might have been the Cardinal.

Here are two observations about the kind of change in information each of these reports would elicit: first, the effects of updating on (1), (2) and (3) are the same. In response to an assertion of any of the first three sentences, a credulous agent ought to rule out the possibility that the Bishop did it. Second, the effect of updating on (4) is different. In response to an assertion of the latter sentence, a credulous agent ought to rule in the possibility that the Cardinal did it. Put another way, in coming to accept one of the first three sentences an agent can acquire new information; in coming to accept the fourth sentence, however, she can only lose information (cf. Stalnaker (2014, 48)).

This pattern forms the basis of the puzzle we'll be exploring. Where ϕ belongs to the non-modal fragment of the language, the following four conditions offer a generalization of our observations about the specific case above.[1]

Transparency	Updating on ϕ and on $\Box\phi$ have the same effect on an agent's information.
Duality	Updating on $\neg\Diamond\phi$ and on $\Box\neg\phi$ have the same effect on an agent's information.
Strength	Updating on $\Box\phi$ never makes a (coherent) agent's information (strictly) weaker than it was previously.
Weakness	Updating on $\Diamond\phi$ only ever makes an agent's information (non-strictly) weaker than it was previously.

Transparency says that embedding a non-modal sentence under an epistemic necessity modal doesn't changes its effect on an audience's information. In our example, this amounts to the requirement that updating on (1) and (2) changes your information in the same way. **Duality** says that epistemic necessity modals and epistemic possibility modals are duals with respect to their effects on an audience's information. In combination with **Transparency**, this amounts to the requirement that (3) changes your information in the same way as (1) and (2).[2]

[1] We'll get clearer on the syntactic details of the language in question later, in §4. For now, it suffices that ϕ belongs to the non-modal fragment iff it contains no instances of $\Diamond$ or $\Box$.

[2] Assuming, at least, that in the context in which your information leaves open only Abbot and Bishop possibilities, 'It can't have been the Bishop' will have the same effect as 'It can't have not been the Abbot'.

Strength says that updating with an statement of epistemic necessity never merely leads to giving up information (unless you're incoherent to start with). Together with **Duality** and **Transparency**, it implies that, if updating on (1)-(3) has a non-trivial effect, the resulting information will rule out something which was previously ruled in. **Weakness** says that updating with a statement about epistemic possibility never leads to acquiring information. It implies that if updating on (4) has a non-trivial effect, the resulting information will rule in something that was previously ruled out (and will not rule out anything that was previously ruled in). Each of these four descriptive conditions seems independently plausible. When combined, however, they have the additional virtue of predicting exactly the pattern of behavior exhibited by (1)-(4).

Two features of the kind of communication we are considering are worth emphasizing. First, it is crucial to our initial case that it is a case of deference: whatever report the detective makes, it is assumed that you (the audience) will update your information on their assertion accordingly. Second, it is important that we are aiming to characterize only the essential effects of assertion on the audience's information (Stalnaker (1978, 2002)). An assertion's essential effect is the change that an audience member must make to their information to count as having accepted it. This need not—and typically will not—be the totality of information which an audience is able to extract from the fact that it occurred. An assertion of (4), for example, may well put its audience in a position to infer a range of new information about the speaker, such as the kind of evidence she possesses, the language she speaks, and so on. However, communicating this information is not part of its essential effect. An audience member who failed to add to her information that the speaker speaks English would not thereby count as failing to accept the assertion.

The effect of update can characterized in terms of an update rule: a procedure which, given an arbitrary sentence and a prior information state, returns a new information state (representing the effect of learning the former in the latter). The four descriptive conditions, above, impose constraints on a rule's extension. However, we can also consider theoretical constraints on how the rule is calculated. The following, in particular, will play an important role in guiding our investigation.

Content The effect of updating on ϕ is a function of the content of ϕ and the agent's prior information.

Content says that update is semantic—to identify the effect of a given update, we need only consider an agent's prior information and the content of the expression she is updating on. This ensures that any two sentences with the same content will have the same effect on an agent's information; it excludes, for example, the possibility that the effect of a sentence on an agent's information can vary simply according to whether or not it belongs to the modal or non-modal fragment of the language.

The strongest argument in favor of **Content** is that it is one of the core components of a simple theory of communication for non-modal languages. According to this theory, agents updating on a utterance simply make the minimal modification to their information necessary to ensure that it incorporates the content of the sentence uttered (Lewis (1969); Stalnaker (1978); Bach & Harnish (1979); Gardenfors (1988)). Update on an utterance can therefore be subsumed under the more general activity of information change, with the content of the sentence uttered playing the role of incoming information.

This picture of communication is elegant, intuitive and widely endorsed. As such, it would be nice to have a theory which extends it to epistemic modality. The problem is that, on some of the most common approaches to theorizing about epistemic modals, it is in tension with our initial observations about communication with epistemic modals. In fact, as I show in §§3-4, both dynamic and propositional theories are incapable of satisfying the combined packed of **Transparency**, **Duality**, **Strength**, **Weakness** and **Content** without lapsing into triviality.

Of course, none of these conditions are beyond dispute. Any of them could reasonably be rejected were the costs of satisfying it shown to be sufficiently great. One response to the results below would be to argue that we should set our sights lower—we should not expect a successful theory of modal communication to satisfy all five conditions. Indeed, nothing that we have said so far has established that they are even mutually consistent.

My primary aim, in what follows, is to argue that we can be more ambitious. I'll show how a theory which satisfies all the conditions above can be constructed. This theory offers a generalization of AGM revision to non-propositional theories of epistemic modality (such as Veltman (1985); Yalcin (2007, 2011); Hawke & Steinert-Threlkeld (2018, 2020), and Rothschild & Yablo (2021)). An important upshot of this is that we can, if we want, extend the simple picture of communication-as-information-change to epistemic modality. The paper concludes with a representation theorem (**Theorem 1**), which establishes the consistency of the theory and shows how it can be understood intuitively in terms of a sequence of fallback information states an agent could occupy.

3 Taxonomy

Theories of modal communication come in two kinds: dynamic theories and propositional theories. The two differ in their division of labour between semantic content and update rules. According to dynamic theories, the content of a sentence itself determines an update on information. According to static theories, in contrast, an update rule must be specified independently of sentences' contents.

3.1 Dynamic Theories

On dynamic theories of modal communication, the content of a sentence is an operation on information. As Groenendijk & Stokhof (1991a) put it:

> "The meaning of a sentence does not lie in its truth conditions, but rather in the way it changes [...] the information of the interpreter"(43)

The guiding idea behind dynamic theories is that instructions on how an agent is to change her information in response to an utterance are encoded directly in the content of the sentence uttered. Accordingly, such theories can offer an extremely simple update rule:

Dynamic Update The effect of updating on ϕ is the result of applying the content of ϕ to the agent's prior information.

Dynamic theories face a choice about the type of operations they treat as the content of sentences. In update semantics (Veltman (1996)),[3] these operations are assumed to be deterministic. Applying the same operation to the same input always returns the same result. Nothing in our characterization of dynamic theories requires this, however (see Hoare (1969)). Willer (2018) proposes permitting non-deterministic operations as the content of sentences in order to model free choice effects. In this framework, the content of a disjunction may return more than one output, given the same input.

Any deterministic dynamic theory which adopts **Dynamic Update** will evidently satisfy **Content**. However, existing dynamic theories do less well with the four descriptive conditions. Update semantics fails to satisfy both **Transparency** and **Weakness**. The same holds of more sophisticated theories, such as Willer (2013) or Willer (2018).As we will see later (§4), this is no coincidence. In fact, no non-trivial dynamic theory is capable of satisfying all five conditions.

3.2 Static Theories

According to static theories, the content of a sentence is not an operation on information. As such, the associated update rules take on a more significant role. Static theories come in two kinds: propositional theories and non-propositional theories.

3.2.1 Propositional Static Theories

On propositional theories, the content of a sentence is a proposition. Propositions are information (Stalnaker (1984, 1999)). They are the kind of objects which rule some ways things could be in and rule other ways things could be out. In this way, propositional theories represent sentential contents as objects

[3]See also Groenendijk & Stokhof (1991b); Groenendijk *et al.* (1996); van der Does *et al.* (1997); Beaver (2001); Gillies (2001, 2020); Willer (2015).

of the same type as an agent's information state. The approach to update these theories adopt varies according to how they go on to characterize those objects.

On the meta-semantic approach defended by Lewis (1979b) and Stalnaker (2014), propositions are identified with properties of worlds (or, less neutrally, with sets of worlds). They are objects which take a stance on the way the world is (and on nothing else). Sentences in the modal fragment of the language are assumed to be context-sensitive: which proposition they express can vary depending on the information possessed by the speaker (Kratzer (1977, 1981, 2012)).

According to proponents of the meta-semantic approach, modal communication proceeds indirectly. When an agent updates on a modal statement like (2) or (4), she does not simply add its content to her information (as she would for a non-modal sentence like (1)). Instead, she does something a bit more sophisticated: she modifies her own information to make it the case that, were she to utter the same sentence, its contextually determined content would be true. For example, in updating on the statement that it might have been the Cardinal, an agent who entertains only the Abbot and the Bishop as suspects will have to rule in some other ways the world might be.

Meta-semantic approaches to a propositional static theory are capable of accommodating all four descriptive conditions. However, **Content** fails, since update on a modal sentence is sensitive, not to the content the sentence in fact has, but rather to the content it would have in different counterfactual contexts. Accordingly, updating on modal and non-modal sentences with the same content can have different effects.

The primary alternative is a relativist approach (defended by Egan *et al.* (2005); Egan (2007, 2018); Stephenson (2007a,b)). Relativists agree that objects of the same type—propositions—play the role both of sentential contents and agents' information states. However, they identify propositions, not with sets of worlds, but with properties of individuals (or, less neutrally, sets of centered worlds Lewis (1979a)).They are objects which take a stance not only on the way the world is, but also the way a person is situated within it. Thus, agents who agree on what the world is like but disagree about what information they possess about it will be represented as endorsing distinct propositions.

Within this richer space of propositions, Egan argues, update can proceed via a single rule (Egan (2007, 2018)). Any differences in the effects of update on modal and non-modal sentences are due to differences in content. In updating on (1), an agent changes her view on how the world is—taking the Abbot to be guilty. In contrast, in updating an (2), she changes her view on the information she possesses—taking herself to be the kind of agent who rules anybody but the Abbot's guilt.

Relativist approaches are designed to satisfy **Content**. However, like dynamic theories, they are incapable of satisfying all four descriptive conditions: **Weakness** fails, since the information of an agent who takes herself to rule out ϕ will not become weaker as a result of updating on $\Diamond\phi$. **Transparency** is also liable to fail. Since ϕ and $\Box\phi$ differ in content, updating on them can be expected to have different effects on an agent's information.

In §4 we will show that, again, this is no coincidence. Any adequate non-propositional static theory must either give up one of the five conditions or make its update rule trivial (in one of two ways).

3.2.2 Non-Propositional Static Theories

On non-propositional theories, the content of a sentence is neither information itself nor an operation on information. The space of possible non-propositional theories is large. However, existing theories have largely coincided in identifying contents with properties of information. That is, contents are the kinds of thing which rule some ways an agent's information could be in and rule others ways it could be out.

Hawke & Steinert-Threlkeld (2018) give a non-propositional theory with an important property: update is sensitive to the syntactic structure of the sentence updated on. Their proposed update rule is designed to satisfy each of our four descriptive conditions. However, in doing so it fails to satisfy **Content**. Updating on two sentences with the same content can have different effects on an agent's information.[4]

An alternative non-propositional static theory is developed in Rothschild & Yablo (2021).[5] Rothschild & Yablo's theory is designed to satisfy **Strength** and **Weakness**. However, their main focus is on issues more fine-grained than those discussed above. In particular, they are interested in explaining which of the weakenings of an agent's information state is returned by updating on $\Diamond\phi$. They achieve this by allowing for failures of **Content**. Specifically, the theory offers separate update rules for $\Diamond$- and $\Box$-embedded sentences, respectively. These rules state the effect on the agent's information not in terms of the content of the sentence itself, but rather the content of its prejacent. Rothschild & Yablo do not intend these rules to extend to update with boolean embeddings of non-modal and modal sentences, since their primary aim is to explain the communicative effects of sentences with modals at widest scope. As they observe, their update rules therefore cover only a partial fragment of the language.

Non-propositional static theories face significant obstacles in satisfying all of our conditions on a theory of modal communication. However, in contrast to dynamic and propostitional static theories, they are not, in principle, incapable of doing so. In §5, I will defend a form of non-propositional theory which meets these requirements. Before turning to this theory, however, it will be helpful to introduce a formal framework for thinking about information change.

[4] For example, $\Diamond\phi \wedge \Diamond(\phi \wedge \psi)$ and $\Diamond(\phi \wedge \psi)$ have the same content according to Hawke & Steinert-Threlkeld (2018). However, they are associated with different update effects.

[5] Rothschild & Yablo (2021) offer their semantics as part of an account of the communicative effects of deontic, not epistemic, modals. However, as previously argued in Yablo (2011), many of the issues raised have the same structure.

4 The Flow of Information

To theorize about information change, we introduce a class of objects: information states. An information state corresponds to a particular view about how things are: it rules some ways things could be out and rules other ways things could be in. Anything that represents things as being a particular way can be represented by an information state. Maps and minds, movies and manuscripts, myths and mime: all belong to kinds of things information states can be used to model. We, however, will employ information states to theorize about just two kinds of things: the contents of sentences and the contents of agents' attitudes.

Information states can be ordered by strength. Where s and s' are information states, $s \leq s'$ iff every way things could be which is ruled out by s' is also ruled out by s. In this case, we say that s is at least as strong as s' or that s incorporates s'. The stronger an information state is, the more comprehensive the view to which it corresponds.

$s \wedge s'$ and $s \vee s'$ are the meet and join of s and s', respectively. $s \wedge s'$ is the weakest state which is stronger than both s and s' (i.e., their unique greatest lower bound). $s \vee s'$ is the strongest state which is weaker than both s and s' (i.e., their unique least upper bound).

An information structure is an ordered set of information states meeting certain conditions (cf. Gardenfors (1988); van Benthem (1986, 1996)).

Definition (Information Structures).

An information structure $\Sigma = \langle \mathbb{S}, \leq \rangle$ is a pair of a countable set of information states $\mathbb{S} = \{s, s', ...\}$ and well-founded partial order, $\leq$, forming a bounded complemented distributive lattice. That is:

- $\mathbb{S}$ has a least element, $\bot$, and greatest element, $\top$.
- Every $s \in \mathbb{S}$ has a unique complement, $\overline{s}$: the state such that $s \wedge \overline{s} = \bot$ and $s \vee \overline{s} = \top$.
- For any $s, s', s'' \in \mathbb{S}$: $s \wedge (s' \vee s'') = (s \wedge s') \vee (s \wedge s'')$.

Intuitively, each information structure corresponds to an exhaustive collection of possible views about how things are. Since it is distributed, bounded, and complemented, each information structure is a boolean algebra. Since $\leq$ is well-founded, it is also a complete boolean algebra. That is, every $S \subseteq \mathbb{S}$ has a unique greatest lower bound, $\bigwedge S$, and a unique least upper bound, $\bigvee S$.

The next thing we need is a language with a modal and a non-modal fragment.

Definition (Language).

- $\mathcal{L}$ is the smallest set containing $\mathcal{A} = \{\mathsf{A}, \mathsf{B}, \mathsf{C}, ...\} \cup \{\mathsf{T}, \mathsf{\bot}\}$ (the set of sentential atoms) which is closed under boolean operators ($\neg$, $\wedge$, $\vee$) and modal operators ($\Diamond, \Box$).

- $\mathcal{L}^-$, the non-modal fragment of $\mathcal{L}$, is the smallest set containing $\mathcal{A}$ which is closed under the boolean operators.

Equipped with the notions of an information structure and a language, we can now compare different theories of how updating on sentences of the language changes an agent's information. In addition to an information structure, each such theory will have two further components: an interpretation of $\mathcal{L}$ (mapping sentences to contents) and an update rule (mapping information states and sentences to information states). In the following two subsections, we will evaluate the prospects of different types of theory satisfying the five conditions.

4.1 Propositional Static Theories

Propositional static theories take the contents of sentences to be information states. A propositional static interpretation function, $[\![\cdot]\!]$, maps sentences of $\mathcal{L}$ to states in an information structure. In a propositional static theory, the boolean constants correspond to boolean operations within an information structures. Negation is identified with complementation (i.e., $[\![\neg\phi]\!]=\overline{[\![\phi]\!]}$), conjunction with meet (i.e., $[\![\phi\wedge\psi]\!] = [\![\phi]\!]\wedge[\![\psi]\!]$) and disjunction with join (i.e., $[\![\phi\vee\psi]\!] = [\![\phi]\!]\vee[\![\psi]\!]$). The tautology and the contradiction denote top and bottom, respectively (i.e., $[\![\mathsf{T}]\!] = \top$ and $[\![\bot]\!] = \bot$). We will not make any assumptions about interpretation of the modal fragment of the language, since this may vary across different propositional theories.

A propositional theory aiming to satisfy **Content** will need to state its update rule in terms of a function which takes a pair of information states as inputs and outputs another information state. That is, it should provide an operation which tells us, if we start in one state, s, and update on a sentence with the content s', what state we will end up in. We can think of such an operation as describing a particular way to move through an information structure.

An appealing thought is that information change should proceed in a minimally disruptive manner (Levi (1977); Makinson (1985)). Revision, under its AGM characterization, is an attempt at implementing this idea (Alchourrón *et al.* (1985)).

Definition (AGM Revision).

$*$ is a revision operation iff $*$ satisfies ($*_1$)-($*_4$):

$$
\begin{array}{lll}
(*_1) & s * s' \leq s' & \\
(*_2) & \bot < s * s' & \textit{if } s' \neq \bot; \\
(*_3) & s * s' = s \wedge s' & \textit{if } s \wedge s' \neq \bot; \\
(*_4) & s * (s' \wedge s'') = (s * s') \wedge s'' & \textit{if } (s * s') \wedge s'' \neq \bot.
\end{array}
$$

AGM revision provides an attractively simple update rule for non-modal communication. In updating on $\phi \in \mathcal{L}^-$, an agent revises her current information with its content; if her old information was s, her new information will be $s * [\![\phi]\!]$.

It is easy to see, however, that this update rule will not extend to modal communication. Consider some non-modal ϕ such that $\Diamond\phi$ has some non-extremal state, s, as its content. That is, $\bot < s < \top$.Then take any s' such that $s \wedge s' < s'$. By $(*_1)$ and $(*_3)$, we know that $s' \not\leq s' * s$; that is, revising s' with the content of $\Diamond\phi$ returns a state either strictly stronger than or incomparable with s'. But **Weakness** says that updating any information state on $\Diamond\phi$ should return a state at least as weak. So **Weakness** will fail.

In fact, we can prove something more general. Say that a theory is successful iff the result of updating on ϕ incorporates $[\![\phi]\!]$. Update via AGM revision is successful (by $(*_1)$). But many weaker revision operations will suffice to guarantee a theory is successful, too.[6] Say that a theory is boring iff, for any $\phi \in \mathcal{L}^-$ and any non-absurd state, the result of updating on ϕ in that state is the same the result of updating either on T or on $\perp$. Boring theories hold that as long as you are in a non-absurd state, any non-modal sentence you update on will have the same effect as learning something trivial or inconsistent. Given these definitions, we can identify an important limitation of propositional static theories (proof in appendix, **§I**).

> **Fact 1.** Any successful propositional static theory which satisfies all five conditions is boring.

Fact 1 says that any propositional static theory meeting the rest of our desiderata must choose between being boring and failing to be successful. Boring theories effectively trivialize the notion of update. On the other hand, any propositional static theory which fails to be successful must deny either (i) that update with ϕ returns an information state at which ϕ is accepted or (ii) that ϕ is accepted at an information state only if that information state incorporates the content of ϕ. Accordingly, it is worth looking to see whether other approaches can do any better.

4.2 Dynamic Theories

Dynamic theories take the contents of sentences to be operations on information states. A dynamic interpretation function, $[\cdot]$, maps sentences of $\mathcal{L}$ to functions in $\mathbb{S} \to \mathbb{S}$ (for a designated information structure). Update proceeds according to **Dynamic Update**: the result of updating on ϕ while in s is $s[\phi]$ (where, following convention, $s[\phi]$ is $[\phi]$ applied to s). We will assume that every dynamic interpretation function associates the tautology with the identity operation on states. That is, for all s, $s[\mathsf{T}] = s$.

For a concrete example, consider the most familiar dynamic theory for $\mathcal{L}$: update semantics (Veltman (1996)).Where v is a function from $\mathcal{A}$ into a designated

[6] Theories strictly weaker than AGM which nevertheless guarantee that $s \odot s' \leq s'$ include Shear & Fitelson (2019)'s lockean theory of revision, Lin & Kelly (2012)'s stability theory and Goodman & Salow (2021, 2023b,a, ms)'s normality-based theory.

information structure Σ, such that $v(\mathsf{T}) = \top$ and $v(\mathsf{\perp}) = \perp$:

Definition (Update Semantics).

- $s[\mathsf{A}] = s \wedge v(\mathsf{A})$
- $s[\neg\phi] = s \wedge \overline{s[\phi]}$
- $s[\phi \vee \psi] = s[\phi] \vee s[\psi]$
- $s[\phi \wedge \psi] = s[\phi][\psi]$
- $s[\Diamond\phi] = \begin{cases} s & \textit{if } s[\phi] \neq \perp; \\ \perp & \textit{otherwise.} \end{cases}$
- $s[\Box\phi] = \begin{cases} s & \textit{if } s[\phi] = s; \\ \perp & \textit{otherwise.} \end{cases}$

Update semantics, paired with an appropriate notion of consequence, generates a non-classical logic in which non-contradicition and excluded middle fail (Mandelkern (2020)). However, it does retain some classical properties. In particular, its interpretation of the boolean operators is congruent: semantic equivalence is invariant under negation, conjunction and disjunction (cf. Humberstone (1995, 2001)). For any $\phi, \psi \in \mathcal{L}$: (i) $[\phi] = [\psi]$ iff $[\neg\phi] = [\neg\psi]$; and (ii) $[\phi]$=$[\psi]$ iff for all $\chi \in \mathcal{L}$, $[\phi \vee \chi] = [\psi \vee \chi]$ and $[\phi \wedge \chi] = [\psi \wedge \chi]$. We will say that a theory is congruential iff its interpretation of the boolean operators is congruent.

It is easy to see that update semantics does not satisfy all four of our descriptive conditions. In particular, **Transparency** and **Weakness** both fail. Where $\phi \in \mathcal{L}^-$, update on ϕ and $\Box\phi$ can have different effects on an agent's information state. And update on $\Diamond\phi$ can take an agent to a strictly stronger information state.[7] In fact we can again prove something more general (building on observations in Rothschild & Yablo (2021, §12)). **Fact 2** identifies an important limitation on dynamic theories (proof in appendix, **§I**).

Fact 2. Any congruential dynamic theory which satisfies all five conditions is boring.

I argued above that a theory must avoid being boring on pain of trivializing update. Should we expect a dynamic theory to be congruential, as well?

Congruence amounts to the requirement that boolean operations preserve sameness and difference in semantic content. Although this is an important feature of classicality, dynamic theorists might, under pressure, learn to live with its failure. However, giving up congruence is insufficient by itself. Any dynamic theory on which negation is compositional (i.e., if $[\phi] = [\psi]$ then $[\neg\phi] = [\neg\psi]$) and which retains equivalence under double negation (i.e., $[\phi] = [\neg\neg\phi]$) will be unable to meet all five conditions without being boring. Giving up either of these properties would be quite radical. Worse, it would be radical while also unmotivated. No data involving epistemic modals suggest that negation might be non-compositional or we should embrace failures equivalence under double negation. Accordingly, any dynamic theory will have to accept substantial costs of some kind.

[7]Specifically, if $\perp < s[\phi] < s$, then $s[\Box\phi] \neq s[\phi]$ and if $\perp = s[\phi] < s$, then $s[\Diamond\phi] < s$.

5 Modal Revision

Both propositional static and dynamic theories face significant limitations. We have seen that, to satisfy our five conditions, they must choose between trivializing their update rule or imposing undesirable constraints on their interpretation functions. A theory which avoids these problems should be at a notable theoretical advantage.

In this section, I show how a theory can be developed which satisfies each of the five conditions without incurring the same costs. At a high level of abstraction, this theory has two component parts: a semantics yielding a particular pattern of acceptance (§5.1) and mechanism for information change based around a series of 'fallback' information states (§5.2). As I show, any theory whose semantics and update rule has these properties will be in a position to satisfy our five conditions.

5.1 Semantics

The theory starts with the idea that whether an agent accepts a sentence is determined by its content and by what information she has. A semantics, supplemented with an acceptance relation, tells us at what states a sentence is accepted. We will focus, in particular, on a family of semantics for $\mathcal{L}$ which exhibit a common set of acceptance properties. Call such semantics 'informational'.

Definition (Informational Semantics).

A semantics (and associated acceptance relation) is informational iff, for any $\phi \in \mathcal{L}^-$:

Downward Persistence	If $\Box\phi$ is accepted at an information state, then $\Box\phi$ is accepted at every information state at least as strong.
Upward Persistence	If $\Diamond\phi$ is accepted at an information state, then $\Diamond\phi$ is accepted at every information state at least as weak.
Synchronic Transparency	$\Box\phi$ is accepted at an information state iff ϕ is accepted at that information state.
Synchronic Duality	$\Box\neg\phi$ is accepted at an information state iff $\neg\Diamond\phi$ is accepted at that information state.

Downward Persistence says that the set of states at which $\Box\phi$ is accepted is downward closed: if it contains some state, s, it also contains any state which rules out every way things could be which is ruled out by s. This reflects the idea that someone accepts $\Box\phi$ iff her information state rules out any way things could be in which ϕ is fails to be the case. Together with **Synchronic Transparency/Duality**, the requirement implies that the set of states at which ϕ and $\neg\Diamond\phi$ are accepted will also be downward closed. **Upward Persistence**

says that the set of states at which $\Diamond\phi$ is accepted is upward closed: if it contains some state s, it also contains any information state which rules in any way things could be which is ruled in by s. This reflects the idea that someone accepts $\Diamond\phi$ iff her information state rules in some way things could be in which ϕ is the case.

Many non-propositional static semantics, supplemented with the appropriate acceptance relation, are informational. For example, Yalcin (2007)'s domain semantics exhibits all four properties, as does Hawke & Steinert-Threlkeld (2018, 2020)'s acceptance semantics and Santorio (2022)'s path semantics. Many dynamic semantics (including update semantics) are also informational semantics when supplemented with a definition of acceptance in terms of fixed pointhood (i.e. s accepts ϕ iff $s[\phi] = s$). Thus, the problems faced by dynamic theories are not, in the end, due to their account of content, but rather due to their adoption of **Dynamic Update** as an update rule.

An informational semantics, by itself, is insufficient. None of the informational systems above offer a theory of communication which satisfies our five conditions. In the next subsection we'll see how to construct an update rule which, combined with any informational semantics on which acceptance is a function of content, yields an adequate theory of modal communication. The generality of this solution is important, since it establishes that the theory's success is not dependent on a specific choice of semantics, but instead on the properties of the update rule.

5.2 Update

Our account of update will start with a simple idea, one familiar from the literature on revision (Levi (1977, 1980); Gärdenfors (1984); Gardenfors & Makinson (1988); Grove (1988)). How an agent responds to incoming information depends on how her prior information is structured. Different parts of her information may be entrenched to different extents and she may, correspondingly, be more readily disposed to give up some parts than others.

We can represent differences in how entrenched parts of an agent's information are using a 'fallback' order. For an agent in the information state s, her associated fallback order is a set of states, $\sigma(s)$, such that: (i) s is the strongest state in $\sigma(s)$; and (ii) for any $s', s'' \in \sigma(s)$, either $s' \leq s''$ or $s'' \leq s'$. The idea is that each state in the set represents the view about how things are that the agent would adopt if forced to give up the information represented by the stronger states in the set. Intuitively, a fallback order can be thought of as akin to a system of spheres in the frameworks of Lewis (1973) and Grove (1988), with information states in the fallback order occupying to role of individual spheres in a sphere system. As an example, figure 1 depicts an information structure and fallback order $\sigma(s_1) = \{s_1, s_2, s_4, \top\}$. States in $\sigma(s_1)$ are shaded blue and connected by dashed edges.

Here's a rough outline of an update rule: in updating on ϕ, an agent starts

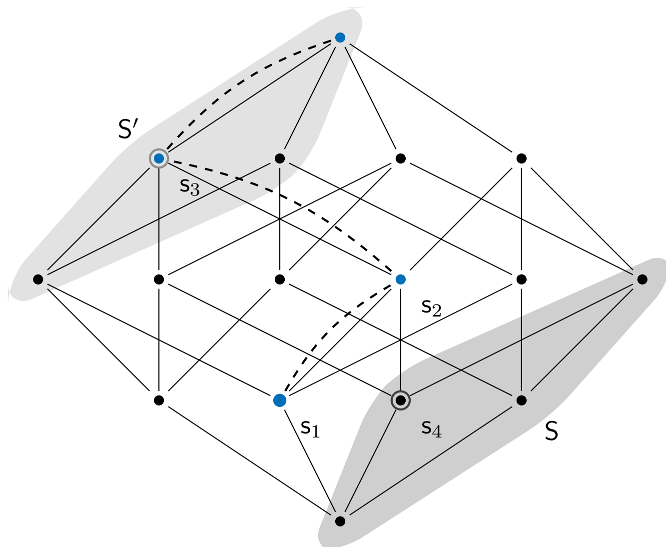

Figure 1: An information structure and fallback order.

by finding the strongest state in her fallback order which can be strengthened into a non-absurd state at which ϕ is accepted. Call this state s'. The agent's new information state will be chosen from among the weakest states at which ϕ is accepted which are at least as strong as s'. This process can be thought of as implementing a form of minimal change: relative to her fallback order, the agent gives up as little old information as possible and acquires as little new information as possible.

As an example, consider the information structure in figure 1. Let S be the downward closed set of information states in the region shaded dark gray. In updating on a sentence accepted at (all and only) the states in this region, an agent who is initially in s_1 must first find the strongest state in the associated fallback order which can be strengthened into a non-absurd state in S. That state is s_2. Her new state will then be among the weakest states in S which are at least as strong as s_2. The unique such state in figure 1 is s_4. In contrast, consider the effect of updating on a sentence accepted (at all and only) the states in S' (the upward closed set of information states in the region shaded light gray). First, the agent will find the strongest state in her fallback order which can be strengthened into a state in S'. That state is s_3. But not only is s_3 the strongest such state, it is also the (unique) weakest state in S' which is at least as strong as s_3. So, her new state will just be s_3 itself.

Even with only an informal (and incomplete) outline, it is possible to see why characterizing update in terms of a fallback order provides a path to satisfying our four descriptive conditions.

Whatever information state an agent is in, the first state in her fallback order is her current state. But recall that, for any non-modal ϕ, the set of states at which $\Box\phi$ is accepted on an informational semantics is downward closed. Suppose that $\Box\phi$ is accepted at some state strictly weaker than the agent's current state. Then it will also be accepted at her current state. So the strongest state in her fallback order which can be strengthened into a state at which $\Box\phi$ is accepted will be her current state itself. And so updating with $\Box\phi$ will never result in a state strictly weaker than an agent's current state, ensuring **Strength** will be satisfied.

Next, recall that for any non-modal ϕ, the set of states at which $\Diamond\phi$ is accepted on an informational semantics is upward closed. Suppose that s' is the strongest state which can be strengthened into a state at which $\Diamond\phi$ is accepted. Then $\Diamond\phi$ will also be accepted at s'. So s' will be the unique weakest state at least as strong as s' at which $\Diamond\phi$ is accepted. And so updating with $\Diamond\phi$ will only ever result in a state (non-strictly) weaker than an agent's current state, ensuring **Weakness** will be satisfied. Finally, on any informational semantics, ϕ and $\Box\phi$ are accepted at the same states, as are $\Box\neg\phi$ and $\neg\Diamond\phi$. So, if the rule treats expressions accepted at the same states alike, **Transparency** and **Duality** will be satisfied.

This is all good. However, this outline falls short of uniquely characterizing an update rule. The procedure described is not (yet) deterministic, since it says only that the result of updating should belong to some (potentially non-singleton) set. Nor is it complete, since it says nothing about what happens when there is no non-absurd strengthening of a fallback state at which the sentence is accepted.

With this in mind, we define a family of 'hyper-revision' operations which map an information state and a property of information states to a new information state. We can think of hyper-revision as a generalization of the notion of revision discussed in §4.1. Revision models the effect of minimally changing one information state so that it incorporates a second, specified, information state. Hyper-revision, in contrast, models the effect of minimally changing an information state so that has some specified property (or, equivalently, so that it belongs to a specified set of information states).

Before we define hyper-revision, some useful notation: ${\downarrow}s = \{s'|s' \leq s\}$ is the set of states at least as strong as s. ${\uparrow}s = \{s'|s' \geq s\}$ is the set of states at least as weak as s. Derivatively, ${\downarrow}S = \bigcup_{s \in S} {\downarrow}s$ and ${\uparrow}S = \bigcup_{s \in S} {\uparrow}s$. $\overline{S} = \mathbb{S}/S$ is the complement of S in the information structure.

Definition (Hyper-revision).

$\circledast$ is a hyper-revision operation iff $\circledast$ satisfies ($\circledast_1$)-($\circledast_7$):

$(\circledast_1)$	$s \circledast S \in S \cup \{\bot\}$		SUCCESS
$(\circledast_2)$	$s \circledast S > \bot$	*if* $S \not\subseteq \{\bot\}$	CONSISTENCY
$(\circledast_3)$	$s \circledast S = s$	*if* $s \in S$ *and* $s > \bot$	TRIVIALITY
$(\circledast_4)$	$s \circledast S \cap S' = s \circledast S$	*if* $s \circledast S \in S \cap S'$	LOCALITY
$(\circledast_5)$	$s \circledast S \cap S' < s' \leq s \circledast S$	*only if* $s' \notin S \cap S'$	MINIMALITY
$(\circledast_6)$	$s \circledast S' \leq s \circledast {\uparrow}S$	*if* $S' \cap {\downarrow}(s \circledast {\uparrow}S) \not\subseteq \{\bot\}$	BOUNDEDNESS
$(\circledast_7)$	$s \circledast S \cap S' = (s \circledast S) \circledast S'$	*if* ${\downarrow}(s \circledast {\uparrow}S) \cap (S \cap S') = {\downarrow}(s \circledast S) \cap S' \not\subseteq \{\bot\}$	DECOMPOSITION

Each of these conditions can be given an independent gloss, in terms of how information should be expected to change in the course of coming to instantiate some property. The first condition says that hyper-revision should be successful: hyper-revising with a set of information states should take you to a state in that set (or to the absurd state). The second condition says that hyper-revision should be consistent: hyper-revising with a set of information states should return a non-absurd state whenever the set contains a non-absurd state itself. Together, the two conditions imply that $s \circledast S = \bot$ iff S is either empty or singleton $\bot$. The third condition says that hyper-revision should not make unnecessary changes: if s is among the states in S, then hyper-revising with S while in s should return s itself, unless s is absurd. The fourth condition says that hyper-revision should be systematic: hyper-revising with any subset of S which contains the result of hyper-revising with S should have the same effect as hyper-revising with S. The fifth condition says that hyper-revision should not introduce more information than necessary: if hyper-revising with a subset of S returns a state at least as strong as the result of hyper-revising with S, then it should be amongst the weakest states in that subset which are at least as strong as the result of hyper-revising with S.

The sixth condition is perhaps less immediately obvious—however, it is crucial to ensuring that hyper-revision implements the intuitive idea of a fallback order. Let s' be the result of hyper-revising with some upward closed set while in s. The condition says that s' must constitute an upper bound on other hyper-revisions of s in a particular way: hyper-revising with any set containing non-absurd states at least as strong as s' will result in a state at least as strong as s'. Informally, the idea is that in hyper-revising with an upward closed set, ${\uparrow}S$, an agent will give up information until a fallback state is reached which belongs to ${\uparrow}S$. Accordingly, the condition encodes the proposal, above, that hyper-revising with a set of states should return a strengthening of the first fallback state which can be strengthened into a non-absurd member of that set.

The seventh condition says that hyper-revision should be incremental: hyper-revising with $S \cap S'$ should have the same effect as first hyper-revising with S and then hyper-revising with S', in the limiting case where: (i) the set of states in $S \cap S'$ which are stronger than the result of giving up the information in $\overline{S}$ is the same as set of states in S' which are stronger than the result of hyper-revising with S, and (ii) there is at least one non-absurd state in each set. The seventh condition ensures that, where there are multiple ways of strengthening a fallback state to reach a state with the desired property, these choices are decided in a systematic way.

With our characterization of hyper-revision in hand, we can state the account of update at the core of our theory. Update is defined in terms of hyper-revision and acceptance in the obvious way:

Hyper-Revisional Update	The effect of updating on ϕ is the result of hyper-revising an agent's prior information with the set of states at which ϕ is accepted.

Hyper-Revisional Update serves as the basis for a theory of modal communication meeting all five of our conditions. Given any informational semantics on which acceptance conditions are determined by content (which includes each of the informational semantic theories discussed above), the theory will clearly satisfy **Content**. An important upshot of this is that the theory will model update, not only for modal-embedded sentences, but also arbitrary boolean combinations of modal- and non-modal-sentences. Most significantly, any such theory will also meet all four descriptive conditions.

Fact 3. Given any informational semantics, **Hyper-Revisional Update** implies **Transparency**, **Duality**, **Strength** and **Weakness**.

Transparency and **Duality** follow immediately from the corresponding synchronic properties of informational semantics, along with the fact that hyper-revision is deterministic.

In this setting, **Strength** says that hyper-revising a non-absurd state with the set of states at which $\Box\phi$ is accepted will never return a strictly weaker state (for $\phi \in \mathcal{L}^-$). To see why **Strength** holds, recall that hyper-revising with any set of states returns either a state in that set (if the the set contains a non-absurd state) or $\bot$ (if it does not). But $\bot$ is not strictly weaker than any state. So suppose that $\Box\phi$ is accepted at some non-absurd state. By **Downward Persistence**, we know that if $\Box\phi$ is accepted at some state strictly weaker than s, then it is also accepted at s. And from $(\circledast_3)$, we know that wherever s is non-absurd, hyper-revising with a set of states containing s while in s will return s itself. So hyper-revising with the set of states at which $\Box\phi$ is accepted will never take a non-absurd state to one which is strictly weaker.

Weakness, in this setting, says that hyper-revising with the set of states at which $\Diamond\phi$ is accepted will only ever return a (non-strictly) weaker state (for $\phi \in \mathcal{L}^-$). To see why **Weakness** holds, observe that every state is non-strictly weaker than $\bot$. So suppose that s is non-absurd. Let s' be the result of hyper-revising, while in s, with the set of states at which $\Diamond\phi$ is accepted. By **Upward Persistence**, we know that this set is upward closed. So, by $(\circledast_6)$, it follows that hyper-revising while in s with any set which contains a non-absurd state at least as strong as s' will return a state at least as strong as s'. So consider the result of hyper-revising while in s with $\{s, s'\}$. Again, we know that the resulting state must be at least as strong as s'. But by $(\circledast_3)$, given that s is non-absurd we know that revising it with $\{s, s'\}$ will return s. Putting the two together, it follows that s is at least as strong as s'. So hyper-revising with the set of states at which $\Diamond\phi$ is accepted will only ever return a state at least as weak as one's original state.

Hyper-revision is, in an important sense, a conservative extension of AGM revision. For any information state s, its downset, $\downarrow s$, is the set comprising all and only those states which carry at least as much information as s. Hyper-revising with $\downarrow s$ can be thought of as making the minimal change required to reach a state which incorporates s. **Fact 4** says that, over such cases, hyper-revision agrees with AGM. For any hyper-revision operation, we can find a corresponding AGM revision operation such that hyper-revising on $\downarrow s$ returns the same state as revising with s (proof in appendix, §**II**).

> **Fact 4.** For any hyper-revision operation, $\circledast$, there is some AGM revision operation, $*$, such that for any $s, s' \in \mathbb{S}$: $s \circledast \downarrow s' = s * s'$.

Finally, we can show that hyper-revision implements in a precise manner the informal idea of moving from an old state to a new state via a pair of successive operations of weakening and strengthening. We establish this in two ways. First, hyper-revision satisfies a generalized form of the Levi identity (Levi (1977)). Just as hyper-revision generalizes AGM revision (**Fact 4**), we can identify corresponding operations generalizing AGM contraction and expansion (appendix, §**III**). $\ominus$ and $\oplus$ denote operations of hyper-contraction and hyper-revision, respectively. Hyper-contraction models the effect of giving up information to find a fallback state which has strengthings lacking some specified property. Hyper-expansion models the effect of adding information to find a state which has some specified property. Crucially, hyper-revision with S is equivalent to sequentially hyper-contracting with $\overline{S}$ and then hyper-expanding with S.

Second, hyper-revision describes a (deterministic, complete) version of the fallback based operation outlined above. Say that $\sigma(s)$ is a fallback order associated with s iff (i) s is the minimum element of $\sigma(s)$; and (ii) $\sigma(s) \subseteq \mathbb{S}$ is totally ordered by $\leq$. Say that $\preccurlyeq$ is a refinement of $\leq$ iff (i) $s \leq s'$ implies $s \preccurlyeq s'$; and (ii) $\mathbb{S}$ is totally ordered by $\preccurlyeq$. For any hyper-revision operation, there is a refinement and a mapping of states to associated fallback orders which characterizes that operation. Conversely, any refinement and mapping from states to associated fallback orders will characterize some hyper-revision operation.

These observations are demonstrated by the following representation theorem (appendix, §**III**).

> **Theorem 1.** Given the axiom of choice, the following are equivalent:

(a) $\circledast$ is a hyper-revision operation;

(b) There is some $\ominus$ and $\oplus$ such that:

   (i) $\ominus$ is a hyper-contraction operation;

(ii) $\oplus$ is a hyper-expansion operation;

(iii) For all s, S: $s \circledast S = (s \ominus \overline{S}) \oplus S$.

(c) There is some $\preccurlyeq$ and σ such that:

(i) $\preccurlyeq$ is a refinement of $\leq$;

(ii) σ is a function which maps each state to a fallback orders associated with that state;

(iii) For all s, S: $s \circledast S = Sup_{\preccurlyeq}(S \cap \downarrow \bigwedge\{s' \in \sigma(s) : \downarrow s' \cap S \not\subseteq \{\bot\}\})$

Theorem 1 says that (for some choice of $\ominus$, $\oplus$, σ and $\preccurlyeq$): $s \circledast S = (s \ominus \overline{S}) \oplus S$ is the result of (i) finding the least $\hat{s}_i \in \sigma(s)$ which is weaker than some non-absurd $s' \in S$ and (ii) returning the $\preccurlyeq$-maximal state in S which is stronger than $\hat{s}_i$.[8]

This theorem not only establishes that our seven conditions on hyper-revision are consistent, it also shows exactly how each hyper-revision operation can be characterized either in terms of a sequence of hyper-contraction and hyper-expansion or in terms of a refinement and mapping from states to associated fallback orders.

6 Conclusion

Epistemic modals present a challenge to our simple picture of communication. I have set out to show how this challenge can be met within a non-propositional static setting.

Many loose threads remain. Nothing has been said about how fallback orders are determined. This means that fine-grained questions about which information agents give up when hyper-revising with $\Diamond\phi$ remain unresolved (cf. Rothschild & Yablo (2021)). Equally, no account has been given of sequential updates. Ideally, it would be good to extend the representation theorem to iterated hyper-revision (cf. Boutilier (1993, 1996); Darwiche & Pearl (1997)). Still, in establishing that the simple picture of communication can be extended to capture update on individual sentences involving modals, we have made a significant first step.

[8]In the case where $S \subseteq \{\bot\}$ (and, hence, there is no such least $\hat{s}_i \in \sigma(s)$), the theorem says that $s \circledast S = \bot$.

Appendix.

I. Where $\circ \in \mathbb{S}^{(\mathbb{S}\times\mathcal{L})}$ is a an update rule and $|\cdot|$ an interpretation function, we can restate our five conditions on theories of communication as follows:

Definition (Conditions).

Content There is some $\odot \in \mathbb{S}^{(\mathbb{S}\times\mathbb{S})}$ such that, for all $\phi \in \mathcal{L} : s{\circ}\phi = s{\odot}|\phi|$.

$$\left.\begin{array}{ll} \textbf{Transparency} & s \circ \phi = s \circ \Box\phi \\ \textbf{Duality} & s \circ \neg\Diamond\phi = s \circ \Box\neg\phi \\ \textbf{Strength} & s \not< s \circ \Box\phi \text{ or } s = \bot \\ \textbf{Weakness} & s \leq s \circ \Diamond\phi. \end{array}\right\} \textit{if } \phi \in \mathcal{L}^-$$

We start by proving **Fact 1**. A theory is boring iff for all $s > \bot$ and $\phi \in \mathcal{L}^-$, either $s \circ \phi = s \circ \mathsf{T}$ or $s \circ \phi = s \circ \bot$. A theory is successful iff $s \circ \phi \leq [\![\phi]\!]$.

Consider a successful propositional static theory with an associated interpretation function, $[\![\cdot]\!]$, and information structure, Σ, which satisfies all five conditions. Since the theory satisfies **Content**, there is some $\odot \in \mathbb{S}^{(\mathbb{S}\times\mathbb{S})}$ such that, for all $\phi \in \mathcal{L} : s \circ \phi = s \odot [\![\phi]\!]$.

Suppose, for contradiction, that the theory is not boring. We'll start by showing that for some $\psi \in \mathcal{L}^-$, $[\![\Diamond\psi]\!] \neq \top$. Given that it is not boring, there is some $\phi \in \mathcal{L}^-$ and some $s > \bot$ such that $s \odot [\![\phi]\!] \neq s \odot [\![\bot]\!]$. From **Transparency**, we have that $s \odot [\![\Box\neg\neg\phi]\!] = s \odot [\![\neg\neg\phi]\!]$. Since $[\![\phi]\!] = [\![\neg\neg\phi]\!]$, it follows that $s \odot [\![\Box\neg\neg\phi]\!] \neq s \odot [\![\bot]\!]$. But we know that $s \odot [\![\Box\neg\neg\phi]\!] = s \odot [\![\neg\Diamond\neg\phi]\!]$, by **Duality**. So $[\![\neg\Diamond\neg\phi]\!] \neq [\![\bot]\!] = \bot$. Since $[\![\neg\Diamond\neg\phi]\!] = \overline{[\![\Diamond\neg\phi]\!]}$ and $\bot = \overline{\top}$, it follows immediately that $[\![\Diamond\neg\phi]\!] \neq \top$.

It follows that there is some $s' \in \mathbb{S}$ such that $s' \not\leq [\![\Diamond\neg\phi]\!]$. Since the theory is successful, $s' \odot [\![\Diamond\neg\phi]\!] \leq [\![\Diamond\neg\phi]\!]$. But, by **Weakness**, $s' \leq s' \odot [\![\Diamond\neg\phi]\!]$. So $s' \leq [\![\Diamond\neg\phi]\!]$, after all. Contradiction. $\Box$

Next, we prove **Fact 2**. Consider a congruential dynamic theory with an associated interpretation function, $[\cdot]$, and information structure, Σ, which satisfies the five conditions.

Consider an arbitrary $\phi \in \mathcal{L}^-$. By **Transparency** and **Duality**, we know that $[\neg\phi] = [\Box\neg\phi] = [\neg\Diamond\phi]$. Since the theory is congruential, it follows that $[\phi] = [\Diamond\phi]$. Next, assume, for reductio, that the theory is not boring. We know that $s[\mathsf{T}] = s$, for all s. So there must be some $\phi \in \mathcal{L}^-$ and $s \neq \bot$ such that $s[\phi] \neq s$. We also know that, for all s, $s[\phi] = s[\Box\phi]$ (by **Transparency**) and $s[\Box\phi] \leq s$ (by **Strength**). So it follows that $s[\phi] < s$. Yet, since $[\Diamond\phi] = [\phi]$, $s[\Diamond\phi] < s$, too. But, by **Weakness**, $s \leq s[\Diamond\phi]$. Contradiction. $\Box$.

II. Our next task is to prove **Fact 4**:

Fact 4. For any hyper-revision operation, $\circledast$, there is some AGM revision operation, $*$, such that for any $s, s' \in \mathbb{S}$: $s \circledast {\downarrow}s' = s * s'$.

Proof: Let $\odot : (\mathbb{S} \times \mathbb{S}) \to \mathbb{S}$ be an operation such that $s \odot s' = s \circledast {\downarrow}s'$. We need to prove that $\odot$ is an AGM revision operation. That is, we must show that $\odot$ satisfies $(*_{1-4})$

First note that, for any $s, s' \in \mathbb{S}$, we know that $s \odot s' \in {\downarrow}s'$ (by $(\circledast_1)$) and $s \circledast \downarrow s' = \bot$ only if $s' = \bot$ (by $(\circledast_2)$). So, it follows that $\odot$ satisfies $(*_1)$ and $(*_2)$.

Observe that $\{s'' \in {\downarrow}s' | s'' \leq s\} \not\subseteq \{\bot\}$ iff $s \wedge s' \neq \bot$. Furthermore, we know that if $\{s'' \in {\downarrow}s' | s'' \leq s\} \not\subseteq \{\bot\}$, then $s \circledast {\downarrow}s' = Max\{s'' \in \downarrow s' | s'' \leq s\} = s \wedge s'$ (by $(\circledast_3)$, $(\circledast_5)$ and $(\circledast_6)$). So it follows that $\odot$ satisfies $(*_3)$, too.

Finally, suppose that $(s \odot s') \wedge s'' \neq \bot$. We need to show that $s \odot (s' \wedge s'') = (s \odot s') \wedge s''$. By hypothesis, and the fact that $\odot$ satisfies, $(*_1)$, we can be sure that $s' \neq \bot$. So we know that $s \circledast {\downarrow}s' = s \circledast ({\downarrow}s' - \{\bot\})$ (by $(\circledast_2)$ and $(\circledast_4)$). Observe that $({\downarrow}s' - \{\bot\}) \subseteq {\uparrow}({\downarrow}s' - \{\bot\})$ and let $s_i = s \circledast {\uparrow}({\downarrow}(s') - \{\bot\})$. From $(\circledast_6)$, it follows that $s \circledast {\downarrow}s' \leq s_i$. Furthermore, by $(\circledast_5)$, $s \odot s'$ is amongst the weakest elements of ${\downarrow}s'$ which are stronger than s_i. But $s' \wedge s_i$ is the unique weakest such state. So $s \odot s' = s_i \wedge s'$.

Next, observe that if $(s \odot s') \wedge s'' \neq \bot$, then ${\downarrow}(s' \wedge s'')) \cap {\downarrow}s_i \not\subseteq \{\bot\}$. So $s \circledast {\downarrow}(s' \wedge s'') \leq s_i$, by $(\circledast_6)$. By hypothesis, $s' \wedge s'' \neq \bot$. So we know that $s \circledast {\downarrow}(s' \wedge s'') = s \circledast ({\downarrow}(s' \wedge s'') - \{\bot\})$ (by $(\circledast_2)$ and $(\circledast_4)$). Furthermore, $({\downarrow}(s' \wedge s'') - \{\bot\}) \subseteq {\uparrow}({\downarrow}s' - \{\bot\})$. So $s \odot (s' \wedge s'')$ is amongst the weakest elements of ${\downarrow}(s' \wedge s'')$ which are stronger than s_i. But $s_i \wedge (s' \wedge s'')$ is the unique weakest such state. So $s \odot (s' \wedge s'') = s_i \wedge (s' \wedge s'')$. But $\wedge$ is associative: that is, $s_i \wedge (s' \wedge s'') = (s_i \wedge s') \wedge s''$. So $(s \odot s') \wedge s'' = s \odot (s' \wedge s'')$. So it follows that $\odot$ satisfies $(*_4)$.

III. Before proving the representation theorem, we introduce some definitions.

Definition (Fall-Back Orders and Refinements). Let σ be mapping of states to *associated fallback orders* iff, for all s: $\bigwedge \sigma(s) = s$ and $\sigma(s) \subseteq \mathbb{S}$ is totally ordered by $\leq$.

Let $\preccurlyeq$ be a *refinement* of $\leq$ iff $s \leq s'$ implies $s \preccurlyeq s'$ and $\mathbb{S}$ is totally ordered by $\preccurlyeq$.

Definition (Hyper-contraction and Hyper-expansion).

$\ominus$ is a *hyper-contraction operation* iff $\ominus$ satisfies $(\ominus_1)$-$(\ominus_5)$.

$$\begin{array}{lll}
(\ominus_1) & {\downarrow}(s \ominus S) \cap \overline{S} \not\subseteq \{\bot\} & \textit{if } \overline{S} \not\subseteq \{\bot\} \\
(\ominus_2) & s \ominus S \geq s & \\
(\ominus_3) & s \ominus S = s & \textit{if } s \in {\uparrow}(\overline{S}/\bot) \\
(\ominus_4) & s \ominus S \cap S' = s \ominus S & \textit{if } \overline{S} \cap {\downarrow}(s \ominus S \cap S') \not\subseteq \{\bot\} \\
(\ominus_5) & s \ominus S' \leq s \ominus {\downarrow}S & \textit{if } \overline{S'} \cap {\downarrow}(s \ominus {\downarrow}S) \not\subseteq \{\bot\}
\end{array}$$

$\oplus$ is a *hyper-contraction operation* iff $\oplus$ satisfies $(\oplus_1)$-$(\oplus_5)$.:

$$\begin{array}{lll}
(\oplus_1) & s \oplus S \in S \cup \{\bot\} & \\
(\oplus_2) & s \oplus S = s & \textit{if } s \in S \\
(\oplus_3) & s \oplus S \cap S' = s \oplus S & \textit{if } s \oplus S \in S \cap S' \\
(\oplus_4) & s \oplus S \cap S' < s' \leq s \oplus S & \textit{only if } s' \notin S \cap S' \\
(\oplus_5) & s \oplus S \cap S' = (s \oplus S) \oplus S' & \textit{if } {\downarrow}s \cap (S \cap S') = {\downarrow}(s \oplus S) \cap S' \not\subseteq \{\bot\}
\end{array}$$

To prove **Theorem 1**, we will prove that (a)$\Rightarrow$(b); (b)$\Rightarrow$(c); and (c)$\Rightarrow$(a).

(a)$\Rightarrow$(b): For an arbitrary hyper-revision operation, $\circledast$, let $\ominus$ and $\oplus$ be defined such that:

$$s \ominus S = s \vee (s \circledast {\uparrow}(\overline{S}/\bot))$$

$$s \oplus S = s \circledast ({\downarrow}s) \cap S.$$

To prove that (a)$\Rightarrow$(b), we need to show that:

i. $\ominus$ is a hyper-contraction operation.

ii. $\oplus$ is a hyper-expansion operation.

iii. For all $s, S : s \circledast S = (s \ominus \overline{S}) \oplus S$.

i. We start by proving the following lemma.

Lemma 1. $s \ominus S = \begin{cases} s \circledast {\uparrow}(\overline{S}/\bot) & \textit{if } \overline{S} \not\subseteq \{\bot\}; \\ s & \textit{otherwise.} \end{cases}$

Proof: Let $s' = s \circledast {\uparrow}(\overline{S}/\bot)$. Suppose that $\overline{S} \not\subseteq \{\bot\}$. Then by $(\circledast_2)$, it follows that $\{s, s'\} \cap {\downarrow}s' \not\subseteq \{\bot\}$. So, by $(\circledast_6)$, $s \circledast \{s, s'\} \leq s'$. By $(\circledast_3)$, either $s = \bot$ or $s \circledast \{s, s'\} = s$. Either way, $s \leq s'$. So $s \ominus S = s \vee s' = s' = s \circledast {\uparrow}(\overline{S}/\bot)$.

Suppose, instead, that $\overline{S} \subseteq \{\bot\}$. Then, by $(\circledast_1)$, $s' = (s \circledast \varnothing) = \bot$. So $s' \leq s = s \vee s' = s' = s \ominus S$.

Next, we show that $\ominus$ satisfies $(\ominus_1) - (\ominus_5)$.

By **Lemma 1**, $(\circledast_1)$ and $(\circledast_2)$, $s \ominus S \in {\uparrow}(\overline{S}/\bot)$, if $\overline{S} \not\subseteq \{\bot\}$. It follows that if $\overline{S} \not\subseteq \{\bot\}$ there is some $s' \in \overline{S}/\bot$ such that $s \ominus S \geq s'$. So $\ominus$ satisfies $(\ominus_1)$. Trivially, $s \vee (s \circledast {\uparrow}(\overline{S}/\bot)) \geq s$. So $\ominus$ satisfies $(\ominus_2)$. Assume $s \in (\overline{S}/\bot)$. Then, by $(\circledast_3)$, $s \circledast (\overline{S}/\bot) = s$. It follows that $s \ominus S = s \vee (s \circledast (\overline{S}/\bot)) = s$. So $\ominus$ satisfies $(\ominus_3)$.

Assume that $\overline{S} \cap {\downarrow}(s \ominus S \cap S') \not\subseteq \{\bot\}$. We'll consider two cases. Suppose, first, that $\overline{S \cap S'} \subseteq \{\bot\}$. Since $\overline{S} \subseteq \overline{S \cap S'}$, it follows that $\overline{S} \subseteq \{\bot\}$. So, by **Lemma 1**, $s \ominus S = s \ominus S \cap S' = s$. So suppose, instead, that $\overline{S \cap S'} \not\subseteq \{\bot\}$. Then, by **Lemma 1**, $s \ominus S \cap S' = s \circledast {\uparrow}(\overline{S \cap S'}/\bot)$. But, since $\overline{S} \cap {\downarrow}(s \ominus S \cap S') \not\subseteq \{\bot\}$, it follows that $s \circledast {\uparrow}(\overline{S \cap S'}/\bot) \in {\uparrow}(\overline{S}/\bot)$. So, by $(\circledast_4)$, $s \circledast {\uparrow}(\overline{S}/\bot) = s \circledast {\uparrow}(\overline{S \cap S'}/\bot)$. But, by assumption, $\overline{S} \not\subseteq \{\bot\}$. So, by **Lemma 1**, $s \ominus S = s \circledast {\uparrow}(\overline{S}/\bot)$, and hence $s \ominus S = s \ominus S \cap S'$. So $\ominus$ satisfies $(\ominus_4)$.

Finally, assume that $\overline{S'} \cap {\downarrow}(s \ominus \overline{\uparrow S}) \not\subseteq \{\bot\}$. It follows that ${\uparrow}(\overline{S'}/\bot) \cap {\downarrow}(s \ominus \overline{\uparrow S}) \not\subseteq \{\bot\}$. We'll consider two cases. First, suppose $\overline{\overline{\uparrow S}} \subseteq \{\bot\}$. Then $s \ominus \overline{\uparrow S} = s$. But since $s = s \circledast {\uparrow} s$, it follows from $(\circledast_6)$ that $s \circledast {\uparrow}(\overline{S'}/\bot) \leq s = s \ominus \overline{\uparrow S}$. So $s \vee (s \circledast {\uparrow}(\overline{S'}/\bot)) = s \ominus S' = s = s \ominus \overline{\uparrow S}$. So suppose, instead, that $\overline{\overline{\uparrow S}} \not\subseteq \{\bot\}$. Then $s \ominus \overline{\uparrow S} = s \circledast {\uparrow}(\overline{\overline{\uparrow S}}/\bot) = s \circledast {\uparrow}(S/\bot)$. It follows from $(\circledast_6)$ and our original assumption that $s \circledast {\uparrow}(\overline{S'}/\bot) \leq s \circledast {\uparrow}(S/\bot)$. But, by the above reasoning, we also know $s \leq s \circledast {\uparrow}(S/\bot)$. So $s \vee (s \circledast {\uparrow}(\overline{S'}/\bot)) = s \ominus S' \leq s \circledast {\uparrow}(\overline{S}/\bot) = s \ominus \overline{\uparrow S}$. So $\ominus$ satisfies $(\ominus_5)$.

ii. We show that $\oplus$ satisfies $(\oplus_1) - (\oplus_5)$.

That $\oplus$ satisfies $(\oplus_1)$ follows immediately from $(\circledast_1)$. That $\oplus$ satisfies $(\oplus_2)$ follows from $(\circledast_1)$ and $(\circledast_3)$. That $\oplus$ satisfies $(\oplus_3)$ follows immediately from $(\circledast_4)$. Suppose $s \oplus (S \cap S') < s' \leq s \oplus S$. Then $s' \notin (S \cap S') \cap {\downarrow} s$, by $(\circledast_5)$. But $s' \in {\downarrow}(s \oplus S) \subseteq {\downarrow} s$. So $s' \notin S \cap S'$. So $\oplus$ satisfies $(\oplus_4)$.

Assume that ${\downarrow} s \cap (S \cap S') = {\downarrow}(s \oplus S) \cap S' \neq \varnothing$. It follows that $({\downarrow} s) \cap S \not\subseteq \{\bot\}$. So $s \in {\uparrow}(({\downarrow} s) \cap S))$. Furthermore, since ${\downarrow} s \not\subseteq \{\bot\}$, $s \neq \bot$. So, by $(\circledast_3)$, it follows that ${\downarrow}(s \circledast {\uparrow}(({\downarrow} s) \cap S)) = {\downarrow} s$. Combining this with our assumption, it follows that ${\downarrow}(s \circledast {\uparrow}(({\downarrow} s) \cap S)) \cap (S \cap S') = {\downarrow}(s \circledast ({\downarrow} s) \cap S) \cap S' \not\subseteq \{\bot\}$. So, by $(\circledast_7)$, we have that $(s \circledast ({\downarrow} s) \cap S) \circledast S' = s \circledast ((({\downarrow} s) \cap S) \cap S')$. Or, equivalently, $(s \oplus S) \circledast S' = s \oplus S \cap S'$.

By $(\circledast_3)$, we have that $s \oplus S = (s \oplus S) \circledast {\uparrow}(s \oplus S)$. Since ${\downarrow}(s \oplus S) \cap S' \not\subseteq \{\bot\}$, it follows from $(\circledast_6)$ that $(s \oplus S) \circledast S' \leq s \oplus S = (s \oplus S) \circledast {\uparrow}(s \oplus S)$. So, by $(\circledast_1)$, $(s \oplus S) \circledast S' \in {\downarrow}(s \oplus S) \cap S'$. Accordingly, by $(\circledast_4)$, we can conclude that $(s \oplus S) \circledast ({\downarrow}(s \oplus S) \cap S') = (s \oplus S) \circledast S'$. Putting the steps together, we get that $(s \oplus S) \oplus S' = s \oplus S \cap S'$. So $\oplus$ satisfies $(\oplus_5)$.

iii. Finally, we show that for an arbitrary s and S: $s \circledast S = (s \ominus \overline{S}) \oplus S$.

Using **Lemma 1**, we consider two cases. Suppose $S \subseteq \{\bot\}$. Observe that by $(\circledast_1)$, for any s', $s' \circledast S' = \bot$, if $S' \subseteq \{\bot\}$. So $(s \ominus \overline{S}) \circledast ({\downarrow} s \cap S) = s \circledast S = \bot$. Suppose instead that $S \not\subseteq \{\bot\}$. Then $s \ominus \overline{S} = s \circledast {\uparrow}(S/\bot)$. We're going to use $(\circledast_7)$ to prove that $s \circledast S = (s \circledast {\uparrow}(S/\bot)) \circledast ({\downarrow}(s \circledast {\uparrow}(S/\bot)) \cap S)$. By $(\circledast_1)$, $(\circledast_2)$ and $(\circledast_5)$, if $S \not\subseteq \{\bot\}$, then for all s: $s \circledast (S/\bot) = s \circledast S$. So it will suffice to show that $s \circledast (S/\bot) = (s \circledast {\uparrow}(S/\bot)) \circledast {\downarrow}(s \circledast {\uparrow}(S/\bot)) \cap (S/\bot)$.

Observe that $\uparrow X \cap X = X$. So it follows that that $\downarrow(s \circledast \uparrow\uparrow(S/\bot)) \cap \uparrow(S/\bot) \cap (S/\bot) = \downarrow(s \circledast \uparrow(S/\bot)) \cap (S/\bot)$. Furthermore, since $S \not\subseteq \{\bot\}$, it follows from($\circledast_1$) and ($\circledast_2$) that $\downarrow(s \circledast \uparrow(S/\bot)) \cap (S/\bot) \not\subseteq \{\bot\}$. So the antecedent of ($\circledast_6$) is satisfied (with $S = \uparrow(S/\bot)$ and $S' = (S/\bot)$). Accordingly, by ($\circledast_6$), we have that $s \circledast \uparrow(S/\bot) \cap (S/\bot) = s \circledast (S/\bot) = (s \circledast \uparrow(S/\bot)) \circledast (S/\bot)$.

Finally, we need to show that $(s \circledast \uparrow(S/\bot)) \circledast (S/\bot) = (s \circledast \uparrow(S/\bot)) \circledast (\downarrow(s \circledast \uparrow(S/\bot)) \cap (S/\bot)) = (s \circledast \uparrow(S/\bot)) \oplus (S/\bot)$. Observe that since $S \not\subseteq \{\bot\}$, it follows that $(S/\bot) \cap \downarrow(s \circledast \uparrow(S/\bot)) \not\subseteq \{\bot\}$. By ($\circledast_3$), since $s \circledast \uparrow(S/\bot) \neq \bot$, it follows that $(s \circledast \uparrow(S/\bot)) \circledast \uparrow(S/\bot) = s \circledast \uparrow(S/\bot)$. So, by ($\circledast_6$), we know that $(s \circledast \uparrow(S/\bot)) \circledast (S/\bot) \in \downarrow((s \circledast \uparrow(S/\bot) \circledast \uparrow(S/\bot)) = \downarrow(s \circledast \uparrow(S/\bot))$. But, equally, by ($\circledast_1$), $(s \circledast \uparrow(S/\bot)) \circledast (S/\bot) \in (S/\bot)$. So $(s \circledast \uparrow(S/\bot)) \circledast (S/\bot) \in \downarrow(s \circledast \uparrow(S/\bot)) \cap (S/\bot)$. Accordingly, by ($\circledast_4$), we can conclude that $(s \circledast \uparrow(S/\bot)) \circledast (S/\bot) = (s \circledast \uparrow(S/\bot)) \circledast (\downarrow(s \circledast \uparrow(S/\bot)) \cap (S/\bot))$.

Putting the steps together, we get that $s \circledast (S/\bot) = (s \circledast \uparrow(S/\bot)) \circledast \downarrow(s \circledast \uparrow(S/\bot)) \cap (S/\bot) = (s \ominus \overline{S}) \oplus S$. But, as observed above, if $S \not\subseteq \{\bot\}$, it follows that $s \circledast S = s \circledast (S/\bot)$. So $s \circledast S = (s \ominus \overline{S}) \oplus S$.

(b)⇒(c): For arbitrary $\ominus$ and $\oplus$, let $\sigma : \mathbb{S} \to \mathcal{P}(\mathbb{S})$ and $\preccurlyeq$ be defined as follows.

- $\sigma(s) = \{\hat{s}_0, \hat{s}_1 ...\}$, where $\hat{s}_0, \hat{s}_1, ...$ is a series such that:
    - $\hat{s}_0 = s$;
    - $\hat{s}_n = s \ominus ((\downarrow\hat{s}_{n-1})/\top)$.
- $\check{s}_i \preccurlyeq \check{s}_j$ iff $i \geq j$, where $\check{s}_0, \check{s}_1 ...$, is a series such that:
    - $\check{s}_0 = \top \oplus \mathbb{S}$
    - $\check{s}_n = \top \oplus \mathbb{S}/\{\check{s}_1, ..., \check{s}_{n-1}\}$

To prove (b)⇒(c), we need to show that:

i. σ is a mapping from states to associated fallback orders.

ii. $\preccurlyeq$ is a refinement of $\leq$.

iii. For all s, S: $(s \ominus \overline{S}) \oplus S = Sup_{\preccurlyeq}(S \cap \downarrow(\bigwedge\{s' \in \sigma(s) : \downarrow s' \cap S \not\subseteq \{\bot\}\}))$

i. Trivially, $\bigwedge \sigma(s) = s$. So we need to show that for all s, $\sigma(s)$ is ordered by $\leq$. For the base case, observe that $\hat{s}_0 \leq s \ominus (\downarrow s/\top) = \hat{s}_1$, by ($\ominus_2$), and $\hat{s}_1 \neq \bot$, by ($\ominus_1$). Suppose, for induction, $\hat{s}_{i-1} \leq \hat{s}_i$ and $\hat{s}_i \neq \bot$. So, $\downarrow\hat{s}_{i-1} \subseteq \downarrow\hat{s}_i$. But, by ($\ominus_1$), $\hat{s}_{i+1} = s \ominus (\downarrow\hat{s}_i/\top) \notin (\downarrow\hat{s}_i/\top)$. So, it follows that $\hat{s}_{i+1} \notin (\downarrow\hat{s}_{i-1}/\top)$, either. By ($\ominus_1$), we know that $\hat{s}_{i+1} \neq \bot$. So $\hat{s}_{i+1} \in (\overline{\downarrow\hat{s}_{i-1}/\top} \cap \downarrow(s \ominus \downarrow\hat{s}_i/\top)) \not\subseteq \{\bot\}$. It follows, by ($\circledast_5$), that $\hat{s}_i = (s \ominus \downarrow\hat{s}_{i-1}/\top) \leq (s \ominus \downarrow\hat{s}_i/\top) = \hat{s}_{i+1}$.

ii. First, we prove that $\mathbb{S}$ is totally ordered by $\preccurlyeq$. It suffices to show that for each $s \in \mathbb{S}$, there is some $i \in \mathbb{N}$ such that $s = \check{s}_i$. Let $S_n = \{\check{s}_0, ..., \check{s}_n\}$. For an arbitrary s, consider the chain $X = \{S_j | s \notin S_j\}$, totally ordered by inclusion. Trivially, $\mathbb{S}$ is an upper bound on X. So, by Zorn's lemma, X has a maximal element, S_k. Suppose, for redutio, $\top \oplus \overline{S_k} \neq s$. Then S_k was not maximal in X, since by construction $S_k \subset (\{s \oplus \overline{S_k}\} \cup S_k) = S_{k+1} \in X$. So $s = \check{s}_{k+1} = s \oplus \overline{S_k}$. But s was arbitrary. So $\mathbb{S}$ is totally ordered by $\preccurlyeq$.

Next, we prove that $\leq$ is included in $\preccurlyeq$. Suppose, for reductio, that $s \neq s'$, $s < s'$ but $s \not\preccurlyeq s'$. Since $\prec$ is (weakly) connected, it follows that $s' \prec s$. So there is some i such that $s, s' \notin S_i$ and $\top \ominus (\mathbb{S}/S_i) = s < s' \leq \top \oplus \mathbb{S}$. But, since $(\mathbb{S}/S_i) \subseteq \mathbb{S}$, by $(\oplus_4)$ we have that $s' \notin (\mathbb{S}/S_i)$. Contradiction. So for distinct s and s' : if $s < s'$, then $s \prec s'$. But $\leq$ and $\preccurlyeq$ are both reflexive and antisymmetric. So $\leq \subseteq \preccurlyeq$.

iii. Finally, we prove that for all s and S: $(s \ominus \overline{S}) \oplus S = Sup_{\preccurlyeq}(S \cap {\downarrow} \bigwedge\{s' \in \sigma(s) : {\downarrow}s' \cap S \not\subseteq \{\bot\}\})$. We consider two cases. Suppose $S \subseteq \{\bot\}$. Then $\bigwedge\{s' \in \sigma(s) : {\downarrow}s' \cap S \not\subseteq \{\bot\}\} \cup \{S \subseteq \{\bot\}\} = \bigwedge \varnothing = \top$. Since $S \subseteq \{\bot\}$, $Sup_{\preccurlyeq}(S \cap {\downarrow}\top) = \bot$. But observe that for all s, where $S \subseteq \{\bot\}$, $s \oplus S = \bot$, by $(\oplus_1)$. So $(s \ominus \overline{S}) \oplus S = \bot = Sup_{\preccurlyeq}(S \cap {\downarrow} \bigwedge\{s' \in \sigma(s) : {\downarrow}s' \cap S \not\subseteq \{\bot\}\})$.

So suppose, instead, $S \not\subseteq \{\bot\}$. We proceed in two steps, by showing that for arbitrary s and $S \not\subseteq \{\bot\}$:

i. $s \ominus \overline{S} = \bigwedge\{s' \in \sigma(s) : {\downarrow}s' \cap S \not\subseteq \{\bot\}\}$;

ii. $s \oplus S = Sup_{\preccurlyeq}(S \cap {\downarrow}s)$.

iii.i. Since $S \not\subseteq \{\bot\}$, there is some least i such that $\hat{s}_i \in \sigma(s)$ and $({\downarrow}\hat{s}_i) \cap S \not\subseteq \{\bot\}$. Since $\hat{s}_i$ is least, ${\downarrow}\hat{s}_{i-1} \cap S \subseteq \{\bot\}$. So it follows that $({\downarrow}\hat{s}_{i-1}/\top) \subseteq {\downarrow}\hat{s}_{i-1} \subseteq (\overline{S/\bot})$. Hence, by elementary set theory, $({\downarrow}\hat{s}_{i-1}/\top) = ({\downarrow}\hat{s}_{i-1}/\top) \cap (\overline{S/\bot})$. By construction, $\hat{s}_i = s \ominus (({\downarrow}\hat{s}_{i-1}/\top) \cap (\overline{S/\bot}))$. So $(S/\bot) \cap {\downarrow}(s \ominus ({\downarrow}\hat{s}_{i-1}/\top) \cap (\overline{S/\bot})) \not\subseteq \{\bot\}$. Accordingly, the antecedent of $(\ominus_4)$ is satisfied, with $S = (\overline{S/\bot})$ and $S' = ({\downarrow}\hat{s}_{i-1}/\top)$. It follows that $s \ominus (\overline{S/\bot}) = s \ominus ({\downarrow}\hat{s}_{i-1}/\top) = \hat{s}_i$. Finally, by $(\ominus_1)$, we know that ${\downarrow}(s \ominus \overline{S}) \cap (S/\bot) \not\subseteq \{\bot\}$. So, by a further application of $(\ominus_4)$, $s \ominus \overline{S} = s \ominus \overline{S/\bot} = \hat{s}_i$. But, by hypothesis, $\hat{s}_i = \bigwedge\{s' \in \sigma(s) : {\downarrow}s' \cap S \not\subseteq \{\bot\}\}$. So $s \ominus \overline{S} = \bigwedge\{s' \in \sigma(s) : {\downarrow}s' \cap S \not\subseteq \{\bot\}\}$.

iii.ii. Next, we show $s \oplus S = Sup_{\preccurlyeq}(S \cap {\downarrow}s)$. We start by proving that $\top \oplus S \cap {\downarrow}s = s \oplus S$. We'll consider two cases. Suppose $S \cap {\downarrow}s \subseteq \{\bot\}$. Then $\top \oplus S \cap {\downarrow}s = \bot = s \oplus S \cap {\downarrow}s$, by $(\oplus_1)$.

So suppose, instead, that $S \cap {\downarrow}s \not\subseteq \{\bot\}$. Observe that, by $(\oplus_2)$ and $(\oplus_4)$, $(\top \oplus {\downarrow}s) \not< s \leq (\top \oplus {\downarrow}\top) = \top$. But, by $(\oplus_1)$, $\top \oplus {\downarrow}s \in {\downarrow}s$. So $(\top \oplus {\downarrow}s) = s$. So it follows that $({\downarrow}\top) \cap (S \cap {\downarrow}s) = S \cap {\downarrow}(\top \oplus {\downarrow}s) = S \cap {\downarrow}s \not\subseteq \{\bot\}$. Accordingly, by $(\oplus_4)$ and commuting $\cap$, we have that $\top \oplus (S \cap {\downarrow}s) = (\top \oplus {\downarrow}s) \oplus S = s \oplus S$.

So we need to prove that $\top \oplus (S \cap {\downarrow}s) = Sup_{\preccurlyeq}(S \cap {\downarrow}s)$. As above, let $S_n = \{\check{s}_0, ..., \check{s}_n\}$. Consider the chain $Y = \{S_j | (S \cap {\downarrow}s) \cap S_j = \varnothing\}$, ordered by inclusion. By Zorn's lemma, Y has some maximal element $S_k = \bigcup Y$. Since S_k is maximal, it follows that $(\top \oplus \overline{S_k}) \in S \cap {\downarrow}s$, since otherwise $S_k \cup (\top \oplus \overline{S_k}) = S_{k+1} \in Y$.

Accordingly, it follows that $(\top \oplus \overline{S_k}) = \check{s}_{k+1} = Sup_{\preccurlyeq}(S \cap {\downarrow}s)$. By construction, $(S \cap {\downarrow}s) \subseteq \overline{S_k}$, and, hence, $(S \cap {\downarrow}s) = (S \cap {\downarrow}s) \cap \overline{S_k}$. It follows, by $(\oplus_3)$, that $\top \oplus (S \cap {\downarrow}s) = \top \oplus (S \cap {\downarrow}s) \cap \overline{S_k} = \top \oplus \overline{S_k}$. So $s \oplus S = \top \oplus (S \cap {\downarrow}s) = Sup_{\preccurlyeq}(S \cap {\downarrow}s)$.

Putting together **iii.i-ii**, we can conclude that $(s \oplus \overline{S}) \oplus S = Sup_{\preccurlyeq}(S \cap {\downarrow} \bigwedge\{s' \in \sigma(s) : {\downarrow}s' \cap S \not\subseteq \{\bot\}\})$

(c)⇒ (b): For an arbitrary refinement, $\preccurlyeq$, and mapping from states to associated fall back orders, σ, let $\circledast$ be defined such that:

$$s \circledast S = Sup_{\preccurlyeq}(S \cap {\downarrow}(\bigwedge\{s' \in \sigma(s) : {\downarrow}s' \cap S \not\subseteq \{\bot\}\}))$$

We need to show that $\circledast$ satisfies $(\circledast_1) - (\circledast_7)$. For ease of comprehension, we'll let $f^\sigma(s, S)$ abbreviate ${\downarrow}(\bigwedge\{s' \in \sigma(s) : {\downarrow}s' \cap S \not\subseteq \{\bot\}\})$.

If $S \cap f^\sigma(s, S) = \varnothing$, then $Sup_{\preccurlyeq}(S \cap f^\sigma(s, S)) = \bot$. Conversely, since $\preccurlyeq$ is total, if $S \cap f^\sigma(s, S) \neq \varnothing$, $Sup_{\preccurlyeq}(S \cap f^\sigma(s, S)) \in S$. So $\circledast$ satisfies $(\circledast_1)$.

Suppose $S \not\subseteq \{\bot\}$. Then $\{s' \in \sigma(s) : {\downarrow}s' \cap S \not\subseteq \{\bot\}\} \neq \varnothing$. Since $\sigma(s)$ is a chain ordered by $\leq$ and $\mathbb{S}$ is complete, $\sigma(s)$ satisfies the descending chain condition. So every non-empty subset of $\sigma(s)$ contains its own meet. But, since $f^\sigma(s, S)$ is the downset of an element of $\{s' \in \sigma(s) : {\downarrow}s' \cap S \not\subseteq \{\bot\}\}$, it follows that $(S \cap f^\sigma(s, S)) \not\subseteq \{\bot\}$. Since $\preccurlyeq$ is a refinement of $\leq$, it follows $Sup_{\preccurlyeq}(S \cap f^\sigma(s, S)) > \bot$. So $\circledast$ satisfies $(\circledast_2)$.

Suppose $\bot < s \in S$. Then $f^\sigma(s, S) = {\downarrow}s$. Since $\preccurlyeq$ is a refinement of $\leq$, it follows $Sup_{\preccurlyeq}(S \cap f^\sigma(s, S)) = s$. So $\circledast$ satisfies $(\circledast_3)$.

Suppose $Sup_{\preccurlyeq}(S \cap f^\sigma(s, S)) \in S \cap S'$. Then $S \cap S' \neq \varnothing$. Suppose $S \subseteq \{\bot\}$. Then, since $\varnothing \subset (S \cap S') \subseteq S \subseteq \{\bot\}$, it follows $S = S \cap S' = \{\bot\}$. So $f^\sigma(s, S) = f^\sigma(s, S \cap S') = {\downarrow} \bigwedge \varnothing = {\downarrow} \top$. So suppose, instead, that $S \subseteq \{\bot\}$. Then, since $\bot < Sup_{\preccurlyeq}(S \cap f^\sigma(s, S)) \in S \cap S'$, it follows that $f(s, S) \cap (S \cap S') \not\subseteq \{\bot\}$. So $f^\sigma(s, S \cap S') = f^\sigma(s, S)$. Since $Sup_{\preccurlyeq}(S \cap f^\sigma(s, S)) \in S \cap S'$ and $Sup_{\preccurlyeq}(S \cap f^\sigma(s, S)) \in f^\sigma(s, S)$, it follows $s \circledast S = Sup_{\preccurlyeq}(S \cap f^\sigma(s, S)) = Sup_{\preccurlyeq}(S \cap S' \cap f^\sigma(s, S \cap S')) = s \circledast S \cap S'$. So $\circledast$ satisfies $(\circledast_4)$.

Suppose $Sup_{\preccurlyeq}((S \cap S') \cap f^\sigma(s, S \cap S')) < s' \leq Sup_{\preccurlyeq}(S \cap f^\sigma(s, S))$. We'll show $s' \notin S \cap S'$. Since $\bot \leq Sup_{\preccurlyeq}((S \cap S') \cap f^\sigma(s, S \cap S')) < s' \leq Sup_{\preccurlyeq}(S \cap f^\sigma(s, S))$, it follows $s' \neq \bot$ and $S \not\subseteq \{\bot\}$. Since $\preccurlyeq$ is a refinement of $\leq$, $Sup_{\preccurlyeq}((S \cap S') \cap f^\sigma(s, S \cap S'))) \prec s'$. So $s' \notin (S \cap S') \cap f^\sigma(s, S \cap S)$. Since $s' \leq Sup_{\preccurlyeq}(S \cap f^\sigma(s, S)) \in f^\sigma(s, S)$ and $f^\sigma(s, S)$ is downward closed, $s' \in f^\sigma(s, S)$. But since $S \cap S' \subseteq S$, it follows that $f^\sigma(s, S) \subseteq f^\sigma(s, S \cap S')$. So $s' \notin S \cap S'$. So $\circledast$ satisfies $(\circledast_5)$.

Suppose $S' \cap {\downarrow}(Sup_{\preccurlyeq}({\uparrow}S \cap f^\sigma(s, {\uparrow}S))) \not\subseteq \{\bot\}$. It follows that $S \neq \varnothing$, (since if $S = \varnothing = {\uparrow}\varnothing$, then ${\downarrow}(Sup_{\preccurlyeq}({\uparrow}S \cap f^\sigma(s, {\uparrow}S))) = {\downarrow}Sup_{\preccurlyeq}(\varnothing) = {\downarrow}\{\bot\} \subseteq \{\bot\})$ and $S' \not\subseteq \{\bot\}$. But observe that if $S \neq \varnothing$, then $\bigwedge\{s' \in \sigma(s) : {\downarrow}s' \cap {\uparrow}S \not\subseteq \{\bot\}\} \in {\uparrow}S$. So ${\downarrow}(Sup_{\preccurlyeq}({\uparrow}S \cap f^\sigma(s, {\uparrow}S))) = {\downarrow}(Sup_{\preccurlyeq}(f^\sigma(s, {\uparrow}S))) = f^\sigma(s, {\uparrow}S)$.

Since $S' \cap f^\sigma(s, \uparrow S) \not\subseteq \{\bot\}$, it follows that $\{\bot\} \subset f^\sigma(s, S') \subseteq f^\sigma(s, S)$. So $Sup_\preccurlyeq(S' \cap f^\sigma(s, S')) \leq f^\sigma(s, \uparrow S) = Sup_\preccurlyeq(\uparrow S \cap f^\sigma(s, \uparrow S))$. So $\circledast$ satisfies $(\circledast_6)$.

Finally we'll prove that $\circledast$ satisfies $(\circledast_7)$. Suppose $\downarrow(Sup_\preccurlyeq(\uparrow S \cap f^\sigma(s, \uparrow S))) \cap (S \cap S') = \downarrow(Sup_\preccurlyeq(S \cap f^\sigma(s, S))) \cap S' \not\subseteq \{\bot\}$. Let $s^* = Sup_\preccurlyeq(S \cap f^\sigma(s, S))$. By supposition, we have that $\downarrow(s^*) \cap S' \not\subseteq \{\bot\}$. So s^* is the least element in $\sigma(s^*)$ such that $\downarrow s^* \cap S' \not\subseteq \{\bot\}$. Hence,$f^\sigma(Sup_\preccurlyeq(S \cap f^\sigma(s, S)), S') = \downarrow(Sup_\preccurlyeq(S \cap f^\sigma(s, S)))$. As above, $\downarrow(Sup_\preccurlyeq(\uparrow S \cap f^\sigma(s, \uparrow S))) = f^\sigma(s, S)$, and hence, $f^\sigma(s, S) \cap (S \cap S') \not\subseteq \{\bot\}$.' So, since $S \cap S' \not\subseteq \{\bot\}$, it follows that $f^\sigma(s, S) = f^\sigma(s, S \cap S')$.

So our original supposition simplifies to the identity $f^\sigma(s, S) \cap (S \cap S') = S' \cap \downarrow Sup_\preccurlyeq(S \cap f^\sigma(s, S))$. It follows immediately that $Sup_\preccurlyeq((S \cap S') \cap f^\sigma(s, S)) = Sup_\preccurlyeq(S' \cap \downarrow(Sup_\preccurlyeq(S \cap f^\sigma(s, S))))$.

But, by the above, $s \circledast (S \cap S') = Sup_\preccurlyeq((S \cap S') \cap f^\sigma(s, S \cap S')) = Sup_\preccurlyeq((S \cap S') \cap f^\sigma(s, S))$ And, similarly, $(s \circledast S) \circledast S' = Sup_\preccurlyeq(S' \cap f^\sigma(Sup_\preccurlyeq(S \cap f^\sigma(s, S)), S')) = Sup_\preccurlyeq(S' \cap \downarrow(Sup_\preccurlyeq(S \cap f^\sigma(s, S))))$. So $s \circledast S \cap S' = (s \circledast S) \circledast S'$. So $\circledast$ satisfies $(\circledast_7)$.

References

Alchourrón, Carlos E., Gärdenfors, Peter, & Makinson, David. 1985. On the Logic of Theory Change: Partial Meet Contraction and Revision Functions. *Journal of Symbolic Logic*, **50**(2), 510–530.

Bach, Kent, & Harnish, Robert M. 1979. *Linguistic Communication and Speech Acts*. Cambridge: MIT Press.

Beaver, David. 2001. *Presupposition and Assertion in Dynamic Semantics*. CSLI Publications.

Boutilier, Craig. 1993 (Jan.). Revision Sequences and Nested Conditionals. *Pages 519–525 of: Proceedings of the 13th International Joint Conference on Artificial Intelligence*.

Boutilier, Craig. 1996. Iterated Revision and Minimal Change of Conditional Beliefs. *Journal of Philosophical Logic*, **25**(3), 263–305.

Darwiche, Adnan, & Pearl, Judea. 1997. On the Logic of Iterated Belief Revision. *Artificial Intelligence*, **89**, 1–29.

Egan, Andy. 2007. Epistemic Modals, Relativism and Assertion. *Philosophical Studies*, **133**(1), 1–22.

Egan, Andy. 2018. De Se Pragmatics. *Philosophical Perspectives*, **32**(1), 144–164.

Egan, Andy, Hawthorne, John, & Weatherson, Brian. 2005. Epistemic Modals in Context. *Pages 131–170 of:* Preyer, Gerhard, & Peter, Georg (eds), *Contextualism in Philosophy*. Oxford University Press.

Gärdenfors, Peter. 1984. Epistemic Importance and Minimal Changes of Belief. *Australasian Journal of Philosophy*, **62**(2), 136–157.

Gardenfors, Peter. 1988. *Knowledge in Flux: Modelling the Dynamics of Epistemic States*. Cambridge, MA.: MIT Press.

Gardenfors, Peter, & Makinson, David. 1988 (Mar.). Revisions of Knowledge Systems Using Epistemic Entrenchment. *Pages 83–95 of:* Vardi, M. Y. (ed), *Proceedings of the Second Conference on Theoretical Aspects of Reasoning about Knowledge*.

Gillies, Anthony S. 2001. A New Solution to Moore's Paradox. *Philosophical Studies*, **105**(3), 237–250.

Gillies, Anthony S. 2020. Updating Data Semantics. *Mind*, **129**(513), 1–41.

Goldstein, Simon. 2021. Epistemic Modal Credence. *Philosophers' Imprint*, **21**(26).

Goodman, Jeremy, & Salow, Bernhard. 2021. Knowledge from Probability. *Electronic Proceedings in Theoretical Computer Science*, **335**(June), 171–186.

Goodman, Jeremy, & Salow, Bernhard. 2023a. Belief Revision from Probability. *Electronic Proceedings in Theoretical Computer Science*, **379**(July), 308–317.

Goodman, Jeremy, & Salow, Bernhard. 2023b. Epistemology Normalized. *The Philosophical Review*, **132**(1), 89–145.

Goodman, Jeremy, & Salow, Bernhard. ms. Belief Revision Normalized.

Groenendijk, Jeroen, & Stokhof, Martin. 1991a. Dynamic Predicate Logic. *Linguistics and Philosophy*, **14**(1), 39–100.

Groenendijk, Jeroen, & Stokhof, Martin. 1991b. Two Theories of Dynamic

Semantics. *Pages 55–64 of:* van Eijck, J. (ed), *Logics in AI*. Berlin, Heidelberg: Springer Berlin Heidelberg.
Groenendijk, Jeroen, Stokhof, Martin, & Veltman, Frank. 1996. Coreference and Modality. *Pages 179–216 of:* Lappin, Shalom (ed), *Handbook of Contemporary Semantic Theory*. Blackwell.
Grove, Adam. 1988. Two Modellings for Theory Change. *Journal of Philosophical Logic*, **17**(2), 157–170.
Hawke, Peter, & Steinert-Threlkeld, Shane. 2018. Informational Dynamics of Epistemic Possibility Modals. *Synthese*, **195**(10), 4309–4342.
Hawke, Peter, & Steinert-Threlkeld, Shane. 2020. Semantic Expressivism for Epistemic Modals. *Linguistics and Philosophy*, 1–37.
Hoare, Tony. 1969. An Axiomatic Basis for Computer Programming. *Commun. ACM*, **12**, 576–580.
Humberstone, I. L. 1995. Negation by Iteration. *Theoria*, **61**(1), 1–24.
Humberstone, Lloyd. 2001. The Pleasures of Anticipation: Enriching Intuitionistic Logic. *Journal of Philosophical Logic*, **30**(5), 395–438.
Kratzer, Angelika. 1977. What 'Must' and 'Can' Must and Can Mean. *Linguistics and Philosophy*, **1**(3), 337–355.
Kratzer, Angelika. 1981. The Notional Category of Modality. *Pages 38–74 of:* Eikmeyer, H. J., & Rieser, H. (eds), *Words, Worlds, and Contexts: New Approaches in Word Semantics*, vol. 6. Berlin: de Gruyter.
Kratzer, Angelika. 2012. *Modals and Conditionals*. Oxford: OUP.
Levi, Isaac. 1977. Subjunctives, Dispositions and Chances. *Synthese*, **34**(4), 423–455.
Levi, Isaac. 1980. *The Enterprise of Knowledge: An Essay on Knowledge, Credal Probability, and Chance*. Vol. 92. MIT Press.
Lewis, David. 1973. *Counterfactuals*. Blackwell.
Lewis, David. 1979a. Attitudes de Dicto and de Se. *Philosophical Review*, **88**(4), 513–543.
Lewis, David. 1979b. Scorekeeping in a Language Game. *Journal of Philosophical Logic*, **8**(1), 339–359.
Lewis, David K. 1969. *Convention: A Philosophical Study*. Wiley-Blackwell.
Lin, Hanti, & Kelly, Kevin. 2012. Propositional Reasoning That Tracks Probabilistic Reasoning. *Journal of Philosophical Logic*, **41**(6), 957–981.
Makinson, David. 1985. How to Give It up: A Survey of Some Formal Aspects of the Logic of Theory Change. *Synthese*, **62**(3), 347–363.
Mandelkern, Matthew. 2020. Dynamic Non-Classicality. *Australasian Journal of Philosophy*, **98**(2), 382–392.
Rothschild, Daniel, & Yablo, Stephen. 2021. Permissive Updates. vol. Outstanding Contributions to Logic: Kit Fine.
Santorio, Paolo. 2018. Conditional Excluded Middle in Informational Semantics. *In: Proceedings of the Amsterdam Colloquium*.
Santorio, Paolo. 2022. Path Semantics for Indicative Conditionals. *Mind*, **131**(521), 59–98.
Shear, Ted, & Fitelson, Branden. 2019. Two Approaches to Belief Revision. *Erkenntnis*, **84**(3), 487–518.

Stalnaker, Robert. 1978. Assertion. *Syntax and Semantics (New York Academic Press)*, **9**, 315–332.
Stalnaker, Robert. 1984. *Inquiry*. Cambridge University Press.
Stalnaker, Robert. 2002. Common Ground. *Linguistics and Philosophy*, **25**(5-6), 701–721.
Stalnaker, Robert. 2014. *Context*. Oxford: Oxford University Press.
Stalnaker, Robert C. 1999. *Context and Content: Essays on Intentionality in Speech and Thought*. Oxford University Press UK.
Stephenson, Tamina. 2007a. Judge Dependence, Epistemic Modals, and Predicates of Personal Taste. *Linguistics and Philosophy*, **30**(4), 487–525.
Stephenson, Tamina. 2007b. *Towards a Theory of Subjective Meaning*. Ph.D. thesis, MIT.
van Benthem, J. F. A. K. 1986. *Essays in Logical Semantics*. Springer.
van Benthem, Johan. 1996. *Exploring Logical Dynamics*. USA: Center for the Study of Language and Information.
van der Does, Jaap, Groeneveld, Willem, & Veltman, Frank. 1997. An Update on 'Might'. *Journal of Logic, Language and Information*, **6**(4), 361–380.
Veltman, Frank. 1985. *Logics for Conditionals*. Ph.D. thesis, University of Amsterdam.
Veltman, Frank. 1996. Defaults in Update Semantics. *Journal of Philosophical Logic*, **25**(3), 221–261.
Willer, Malte. 2013. Dynamics of Epistemic Modality. *Philosophical Review*, **122**(1), 45–92.
Willer, Malte. 2015. An Update on Epistemic Modals. *Journal of Philosophical Logic*, **44**(6), 835–849.
Willer, Malte. 2018. Simplifying with Free Choice. *Topoi*, **37**(3), 379–392.
Yablo, Stephen. 2011. Permission and (So-Called Epistemic) Possibility. *Pages 270–295 of:* Egan, Andy, & Weatherson, Brian (eds), *Epistemic Modality*. Oxford: OUP.
Yalcin, Seth. 2007. Epistemic Modals. *Mind*, **116**(464), 983–1026.
Yalcin, Seth. 2011. Nonfactualism About Epistemic Modality. *In:* Egan, Andy, & Weatherson, B. (eds), *Epistemic Modality*. Oxford University Press.