Autodoxastic Conditional Reasoning: The Monotonic Case⋆

Haythem O. Ismail and Aya S. Mahfouz

German University in Cairo
Department of Computer Science
{haythem.ismail,aya.mahfouz}@guc.edu.eg

Abstract. Ramsey's test for conditionals seems to be in conflict with the so-called Thomason conditionals. A Thomason conditional is a conditional in which either the antecedent or the consequent is a statement about the reasoning agent's own beliefs. Several authors have pointed out that resolving the apparent conflict is to be sought by abandoning the belief revision interpretation of the Ramsey test in favor of a suppositional interpretation. We formalize an AGM-style notion of supposition, showing that it is identical to revision for agents who are not autodoxastic—agents who do not reason about their beliefs. We present particular realizations of supposition in terms of revision and identify the relations between the conditionals supposition and revision give rise to.

1 Ramsey, Thomason, and Moore

In this paper, we attempt to consolidate two issues that, apparently, are in tension: the Ramsey test for conditionals and the so-called Thomason conditionals. The Ramsey test grounds the plausibility of conditionals (roughly, sentences of the form "If P then Q") in a process of belief change. In an often quoted excerpt from [1], Robert Stalnaker gives a procedural interpretation of the Ramsey test:

> First, add the antecedent (hypothetically) to your stock of beliefs; second, make whatever adjustments are required to maintain consistency (without modifying the hypothetical belief in the antecedent); finally, consider whether or not the consequent is then true.

Researchers in artificial intelligence (AI) and philosophical logic will immediately recognize the process Stalnaker is referring to as one of belief revision [2]. This is certainly a welcome interpretation of the Ramsey test, since belief revision has been thoroughly investigated both in AI (for example [3, 4]) and

---

philosophy. Unfortunately, interpreting the Ramsey test as implying a process of belief revision has been dealt at least two blows.

First, Peter Gärdenfors [5] proved that the belief revision interpretation of the Ramsey test is inconsistent with a minimal set of harmless demands on a logical theory. We do not address this problem here; but see [6]. Second, a number of authors, since 1980 [7] till the turn of the decade [8–11], have pointed out that the Ramsey test, à la Stalnaker, provides counter-intuitive judgments of some conditionals—the so-called Thomason conditionals [7].

Thomason conditionals are conditionals in which either the antecedent or the consequent are statements about the reasoning agent's own beliefs. In particular, they come in four main forms:[1]

TC1. If ϕ, then I believe ϕ.
TC2. If I believe ϕ, then ϕ.
TC3. If ϕ, then I do not believe ϕ.
TC4. If I do not believe ϕ, then ϕ.

Given the Stalnaker (belief-revision) interpretation of the Ramsey test, one should accept **TC1** and **TC2**, and reject **TC3** and **TC4**; otherwise, one would succumb to accepting a Moore-paradoxical sentence:

M1. ϕ and I do not believe ϕ.
M2. Not ϕ and I believe ϕ.

But these judgments are not always correct, as several examples attest.

(1) ? If Sally is a spy, then I believe that Sally is a spy.

(2) ? If I believe that Sally is a spy, then Sally is a spy.

(3) If Sally were deceiving me, I would believe that she was not deceiving me (because she is so clever).[2]

(4) Even if I were not to believe that Sally is a spy, she would be a spy (my misconceptions do not change the facts).

A number of authors have attempted to reconcile the Ramsey test with such data [8–11]. The main idea is that Stalnaker's interpretation of the Ramsey test is not exactly faithful to Ramsey's real proposal. However, except for Willer [10] who presents a theory within the framework of update semantics, these proposals are largely informal. In this short report, we propose to stay as close as possible to the AGM belief revision tradition [2]. We introduce an AGM-style belief change operator—supposition—that, we claim, adequately accounts

---

[1] Strictly speaking, the agent's own beliefs need not explicitly appear in a conditional to make the same "Thomason-effect", but even in such cases, beliefs of the agents are assumed to establish the relevance of the antecedent to the consequent.

[2] This is a classical example attributed to Richmond Thomason (hence, "Thomason conditionals"), but appears in print in [7].

for Thomason conditionals if taken to be the belief change operator implicit in Ramsey's test. Unlike [10], however, we retain the common epistemic reading of the Ramsey test side by side with the suppositional one. Hence, we distinguish two classes of conditionals, based on these two readings.

2   Autodoxastic Agents

We call an agent that can reason about its beliefs an autodoxastic agent. Why is it paradoxical for an autodoxastic agent to hold beliefs such as **M1** and **M2**? Intuitively, in normal situations (which we assume throughout), an agent cannot hold that ϕ is true and simultaneously fail to believe it; or hold that ϕ is not true and, nevertheless, believe it. But things are not that simple, or otherwise how would one reply to Chalmers and Hájek's remarks [8] about the oddity of unconditionally accepting **TC1** and **TC2**?

The two pairs (**M1**, **M2**) and (**TC1**, **TC2**) both make sense in some contexts of reasoning, and make no sense in others. In some contexts of reasoning—*regular* contexts—the contexts within which the agent normally and usually reasons, what the agent takes to be true is identical to what it believes. In such contexts, **TC1** and **TC2** are totally acceptable, while **M1** and **M2** are impossible. In *non-regular* contexts, the opposite is true: **M1** and **M2** are possible while **TC1** and **TC2** are usually not acceptable.

What are these non-regular contexts? These are hypothetical contexts of reasoning, in which the agent entertains what the world would be like under one or more suppositions. But once the agent is endowed with the capacity for hypothetical reasoning, it may suppose anything about the world, even things about its own beliefs and their incompatibility with how the (hypothetical) world is. Certainly, the agent may entertain worlds in which ϕ is true, but it does not believe it. Hence, in non-regular contexts **M1** and **M2** are not paradoxical at all. In addition, **TC1** and **TC2** are clearly irrational, in general. Thus, unlike in **TC1** and **TC2**, the "if" in **TC3** and **T4** does not mean "if I accommodate ϕ (respectively, $\neg\phi$) in my regular context" (the belief revision, epistemic reading); rather, it means "if I suppose ϕ (respectively, $\neg\phi$) in a non-regular context" (the suppositional reading).

This being said, our task then is to formalize the distinctions between (i) regular and non-regular contexts, (ii) revision and supposition, and (iii) epistemic and suppositional conditionals.[3]

---

[3] We use "autodoxastic" rather than "autoepistemic" not just to be different from Moore [12]. In Moore's autoepistemic logic, only what correspond to our regular contexts are considered. Thus, for an autoepistemic agent, knowledge and belief are equivalent. In our case, since we consider non-regular contexts of reasoning, using the term "autoepistemic" is at least misleading; we stand by "autodoxastic" then.

3   Truth is in the "I" of the "B" Holder

Suppose we have a logical language with an operator $\mathbf{B}$ for belief and a constant $\mathbf{I}$ that denotes the reasoning agent; the agent whose beliefs are represented by sentences of the language. What characterizes a regular context of reasoning? According to our discussion so far, regular contexts of reasoning should admit $\mathbf{TC1}$ and $\mathbf{TC2}$. Thus, we may label a context "regular" if it admits the following $\mathbf{B}$-schema (after Tarski's T-schema).

$$\phi \equiv \mathbf{B}(\mathbf{I}, \phi)$$

But on closer examination, this formula is too strong. For example, the following line of reasoning is possible.

$$\neg\mathbf{B}(\mathbf{I}, \phi) \Longrightarrow \neg\phi \Longrightarrow \mathbf{B}(\mathbf{I}, \neg\phi)$$

Hence, the following would be derivable:

$$\neg\mathbf{B}(\mathbf{I}, \phi) \supset \mathbf{B}(\mathbf{I}, \neg\phi)$$

The agent, thus, holds an opinion about every proposition. Now, we may try to avoid this by abandoning the $\mathbf{B}$-schema and sticking to Moore's notion of autoepistemic models [12], in which regularity is enforced through non-monotonic rules of inference. This, however, will not be our favorite route in this paper. First, we believe that (a variant of) the $\mathbf{B}$-schema would still be derivable. Second, we will be forced to adopt a non-monotonic logic. In this first report on autodoxastic conditional reasoning, we would like to focus on properties of revision and supposition in the monotonic case, without getting entangled into the complications that a non-monotonic treatment is bound to get us into.

Hence, we shall adopt a weakened version of $\mathbf{B}$-schema. In particular, the problematic direction (viz $\phi \supset \mathbf{B}(\mathbf{I}, \phi)$) will be replaced by the weaker formula

$$\phi \supset \overrightarrow{\mathbf{B}}(\mathbf{I}, \phi)$$

where $\overrightarrow{\mathbf{B}}(\mathbf{I}, \phi)$ is $\mathbf{B}(\mathbf{I}, \phi)$ with $\mathbf{B}$ pushed inside till the outermost positive disjunction or existential quantifier.[4] To make this notion precise, we need to introduce a particular language for reasoning about belief.

4 $Log_A\mathbf{B}$

The language we use in the representation of and reasoning about belief is $Log_A\mathbf{B}$ [14, 15], an algebraic logic of belief, consisting only of sorted terms and no formulas. Instead, a special sort of terms is taken to denote propositions. Semantically, the domain of propositions is a complete, non-degenerate, Boolean algebra [16].

---

[4] A similar manoeuver was made by Perlis [13] with the T-schema in order to avoid paradoxes of self-reference.

Thus reifying propositions provides us with high expressivity compared to standard modal approaches to belief ([17, 18], for instance), while avoiding the complications and susceptibility to self-reference paradoxes that plague syntactical approaches [19, 13].

For example, the sentence "John believes that Sally is a spy" may be represented in $Log_A\mathbf{B}$ by the term

$$\mathbf{B}(J, Spy(S))$$

where $Spy(M)$ is a functional term denoting the proposition that Mary is a spy, while $\mathbf{B}(J, Spy(M))$ is a functional term denoting the proposition that John believes that Mary is a spy. Thus, the belief operator is neither a modal operator nor a predicate symbol, but a function symbol $\mathbf{B}$ denoting a function from pairs of agents and propositions to propositions. The sentence "John believes everything that Sally believes" may be easily expressed in $Log_A\mathbf{B}$:

$$\forall x[\mathbf{B}(S, x) \supset \mathbf{B}(J, x)]$$

For limitations of space, we only give a brief presentation of $Log_A\mathbf{B}$ syntax; for more details about the syntax and semantics, see [14, 15]. A $Log_A\mathbf{B}$ language is a set of terms partitioned into two base syntactic sorts, σ_P and σ_I. Intuitively, σ_P is the set of terms denoting propositions and σ_I is the set of terms denoting anything else. A distinguished subset σ_A of σ_I comprises agent-denoting terms. Standard propositional connectives in B are function symbols on σ_P; quantifiers are treated as punctuation symbols that, together with variables and σ_P terms, construct σ_P terms. The belief operator $\mathbf{B}$ is a dyadic function symbol that forms a σ_P term when combined with, respectively, a σ_A and a σ_P term.

For the purpose of this paper, we introduce to $Log_A\mathbf{B}$ a function symbol $\overrightarrow{\mathbf{B}}$, akin to $\mathbf{B}$, and is interpreted in such a way that it forms terms which are logically equivalent to σ_P terms with occurrences of $\mathbf{B}$ distributed and pushed inwards to the outermost occurrences of positive disjunctive terms. More precisely, if $\llbracket \cdot \rrbracket$ is the semantic interpretation function, then the following restrictions on the interpretation of $\overrightarrow{\mathbf{B}}$ apply; $\alpha \in \sigma_A$, $\{\phi, \psi\} \subseteq \sigma_P$, and π is an atomic σ_P term.

- $\llbracket \overrightarrow{\mathbf{B}}(\alpha, \pi) \rrbracket = \llbracket \mathbf{B}(\alpha, \pi) \rrbracket$
- $\llbracket \overrightarrow{\mathbf{B}}(\alpha, \neg\phi) \rrbracket = \llbracket \neg\overrightarrow{\mathbf{B}}(\alpha, \phi) \rrbracket$
- $\llbracket \overrightarrow{\mathbf{B}}(\alpha, \phi \vee \psi) \rrbracket = \llbracket \mathbf{B}(\alpha, \phi \vee \psi) \rrbracket$
- $\llbracket \overrightarrow{\mathbf{B}}(\alpha, \exists x[\phi]) \rrbracket = \llbracket \mathbf{B}(\alpha, \exists x[\phi]) \rrbracket$

$\wedge, \supset, \equiv$, and $\forall$ are defined in the usual way in terms of $\neg, \vee$, and $\exists$.

In the rest of the paper, we assume a $Log_A\mathbf{B}$ language $\mathcal{L}_0$. If $\mathcal{L} \subseteq \mathcal{L}_0$, then $\mathcal{L}^\tau$ denotes the set of terms of sort τ in $\mathcal{L}$. Such a set $\mathcal{L}$ is propositionally closed if it is the closure of some subset of $\mathcal{L}_0$ under Boolean connectives and quantification.

5   Revision

For the purpose of this paper, we assume contexts to be sets of $\mathcal{L}^P$ terms, where $\mathcal{L}$ is a propositionally-closed subset of $\mathcal{L}_0$. An $\mathcal{L}$ context system is a set $\mathbb{C}$ of

$\mathcal{L}$-contexts satisfying certain closure properties; for example, deductive closure. However, nothing much hangs on what these properties exactly are; some of them will be pointed out as we proceed with the discussion.

For an $\mathcal{L}$ context system $\mathbb{C}$, revision is a mapping $\circledast : \mathbb{C} \times \mathcal{L}^P \longrightarrow \mathbb{C}$. The six basic AGM postulates for revising a context $\mathcal{K}$ with a $\mathcal{L}^P$ term ϕ, denoted $\mathcal{K} \circledast \phi$, are [2]:

$(A^\circledast 1)$ Closure: $\mathcal{K} \circledast \phi \in \mathbb{C}$.[5]
$(A^\circledast 2)$ Success: $\phi \in \mathcal{K} \circledast \phi$.
$(A^\circledast 3)$ Consistency: If $\mathrm{Cn}(\mathcal{K} \circledast \phi) = \mathcal{L}^P$ then $\mathrm{Cn}(\{\phi\}) = \mathcal{L}^P$.
$(A^\circledast 4)$ Expansion 1: $\mathrm{Cn}(\mathcal{K} \circledast \phi) \subseteq \mathrm{Cn}(\mathcal{K} \cup \{\phi\})$.
$(A^\circledast 5)$ Expansion 2: If $\neg\phi \notin \mathrm{Cn}(\mathcal{K})$, then $\mathrm{Cn}(\mathcal{K} \cup \{\phi\}) \subseteq \mathrm{Cn}(\mathcal{K} \circledast \phi)$.
$(A^\circledast 6)$ Extensionality: If $\phi \equiv \psi \in \mathrm{Cn}(\varnothing)$, then $\mathrm{Cn}(\mathcal{K} \circledast \phi) = \mathrm{Cn}(\mathcal{K} \circledast \psi)$.

In the above postulates, where $\Gamma \subseteq \mathcal{L}^P$, $\mathrm{Cn}(\Gamma)$ is the set of consequences of Γ—the closure of Γ under a set of $Log_A\mathbf{B}$ rules of inference. In particular, we envision a set of natural deduction rules as in [6]. Throughout the discussion, Cn is taken to be monotonic.

In the rest of the current section, unless otherwise stated, we consider only the case when $\mathcal{L} = \mathcal{L}_0$. The fundamental distinction we make in autodoxastic context systems is that between regular and non-regular contexts.

Definition 1 *For every $\alpha \in \sigma_A$, a context $\mathcal{K} \in \mathbb{C}$ is α-regular if the following reflection principles are in $\mathrm{Cn}(\mathcal{K})$.*

$(\mathbf{R} \uparrow)$ $\forall p \ [p \ \supset \ \overrightarrow{\mathbf{B}}(\alpha, p)]$ *and*
$(\mathbf{R} \downarrow)$ $\forall p \ [\mathbf{B}(\alpha, p) \supset p]$

The ability to quantify over propositions allows us to succinctly represent the reflection principles. Otherwise, in a modal setting for instance, each of $(\mathbf{R}\uparrow)$ and $(\mathbf{R}\downarrow)$ would correspond to an infinite conjunction of instances. In particular, within $Log_A\mathbf{B}$, we can revise with $[(\mathbf{R}\uparrow) \wedge (\mathbf{R}\downarrow)]$, whereas a modal approach will require an infinitary logical treatment.

The following results immediately follow.[6]

Observation 1 *Let $\mathcal{K}$ be an α-regular context.*

1. *$\mathcal{K}$ is α-consistent: $\forall p \ [\mathbf{B}(\alpha, \neg p) \supset \neg\mathbf{B}(\alpha, p)] \in \mathrm{Cn}(\mathcal{K})$*
2. *$\mathcal{K}$ is α-positively-faithful: $\forall p \ [\mathbf{B}(\alpha, \mathbf{B}(\alpha, p)) \supset \mathbf{B}(\alpha, p)] \in \mathrm{Cn}(\mathcal{K})$*
3. *For all $n \in \mathbb{N}$, $\phi \supset \overrightarrow{\mathbf{B}}^n(\alpha, \phi) \in \mathrm{Cn}(\mathcal{K})$, where $\overrightarrow{\mathbf{B}}^1(\alpha, \phi) = \overrightarrow{\mathbf{B}}(\alpha, \phi)$ and $\overrightarrow{\mathbf{B}}^i(\alpha, \phi) = \mathbf{B}(\alpha, \overrightarrow{\mathbf{B}}^{i-1}(\alpha, \phi))$, for $i > 1$.*

---

[5] Thus, $\mathbb{C}$ is closed under revision with terms in $\mathcal{L}^P$.

[6] All proofs are mostly straightforward, and are deferred, for completeness, to an appendix. For limitations of space, proofs of observations are omitted.

Thus, regularity entails consistency, faithfulness, and a kind of positive introspection with respect to $\overrightarrow{\mathbf{B}}$. Most importantly, however, such a context has the property that ϕ and $\mathbf{B}(\alpha, \phi)$ are almost synonymous. This is the essence of subjective belief as discussed in Section 2.

It should be noted that some anomalous inferences are not possible. For example, $\forall p\ [\neg\mathbf{B}(\alpha, p) \supset \mathbf{B}(\alpha, \neg p)]$ is generally not derivable in an α-regular context, whereas it would have been had we decided to replace $\overrightarrow{\mathbf{B}}$ in $(\mathbf{R}\uparrow)$ with $\mathbf{B}$. (See the discussion of Section 2.)

In addition to the standard AGM postulates, the following postulate constrains the effect of revision on the regularity of the revised context. In essence, it requires the adopted revision strategies to protect $(\mathbf{R}\uparrow)$ and $(\mathbf{R}\downarrow)$ as much as possible.

$(A^{\circledast}7)$ Regularity: If $\mathcal{K}$ is α-regular and $\mathcal{K} \circledast \phi$ is not α-regular, then
    $[\neg(\mathbf{R}\uparrow) \vee \neg(\mathbf{R}\downarrow)] \in \mathrm{Cn}(\{\phi\})$.

From $(A^{\circledast}2)$, $(A^{\circledast}3)$, $(A^{\circledast}7)$, and monotonicity of Cn, $(T^{\circledast}1)$ immediately follows:

$(T^{\circledast}1)$ Regularity preservation: If $\mathcal{K}$ is α-regular, then $\mathcal{K} \circledast \phi$ is not α-regular if and only if $[\neg(\mathbf{R}\uparrow) \vee \neg(\mathbf{R}\downarrow)] \in \mathrm{Cn}(\{\phi\})$.

Example 1. Let $\mathcal{K}$ be an α-regular context. Clearly, $\mathcal{K} \circledast [p \wedge \neg\mathbf{B}(\alpha, p)]$ is not regular. Intuitively, if one is to revise their beliefs with an atomic proposition and the fact that they do not believe that very proposition, then one is implicitly abandoning $(\mathbf{R}\uparrow)$. Similarly, $\mathcal{K} \circledast [\neg p \wedge \mathbf{B}(\alpha, p)]$ is not regular. But, this time, the new information is in direct contradiction with $(\mathbf{R}\downarrow)$.   □

6   Supposition

Corresponding to the basic revision postulates, the following is a set of postulates for idealized supposition. There is one supposition operator $\oplus_{\alpha} : \mathbb{C} \times \mathcal{L}^P \longrightarrow \mathbb{C}$ for every $\alpha \in \mathcal{L}^A$, where $\mathbb{C}$ is an $\mathcal{L}$ context system.

$(A^{\oplus_{\alpha}}1)$ Closure: $\mathcal{K} \oplus_{\alpha} \phi \in \mathbb{C}$.
$(A^{\oplus_{\alpha}}2)$ Success: $\phi \in \mathcal{K} \oplus_{\alpha} \phi$.
$(A^{\oplus_{\alpha}}3)$ Consistency: If $\mathrm{Cn}(\mathcal{K} \oplus_{\alpha} \phi) = \mathcal{L}^P$, then $\mathrm{Cn}(\{\phi\}) = \mathcal{L}^P$.
$(A^{\oplus_{\alpha}}4)$ Extensionality: If $\phi \equiv \psi \in \mathrm{Cn}(\varnothing)$, then $\mathrm{Cn}(\mathcal{K} \oplus_{\alpha} \phi) = \mathrm{Cn}(\mathcal{K} \oplus_{\alpha} \psi)$.

The above postulates are identical to their revision counterparts. Only the expansion postulate are missing. In general, supposition causes more damage to a context than revision; it may add more than the supposed proposition to the context and, even when the supposed proposition is consistent with the context, preservation of the old information is not guaranteed.

In case $\mathcal{L} = \mathcal{L}_0$, the following regularity-related postulates constrain the behavior of $\oplus_{\alpha}$.

$(A^{\oplus_\alpha}5)$ Expansion 1: If $\mathcal{K}\cup\{\phi\}$ is not α-regular, then $\mathrm{Cn}(\mathcal{K}\oplus_\alpha\phi) \subseteq \mathrm{Cn}(\mathcal{K}\cup\{\phi\})$.

$(A^{\oplus_\alpha}6)$ Expansion 2: If $\neg\phi \notin \mathrm{Cn}(\mathcal{K})$ and $\mathcal{K} \cup \{\phi\}$ is not α-regular, then
$\mathrm{Cn}(\mathcal{K} \cup \{\phi\}) \subseteq \mathrm{Cn}(\mathcal{K} \oplus_\alpha \phi)$.

$(A^{\oplus_\alpha}7)$ Regularity: If $\mathcal{K} \oplus_\alpha \phi$ is α-regular, then $[(\mathbf{R} \uparrow) \wedge (\mathbf{R} \downarrow)] \in \mathrm{Cn}(\{\phi\})$.

From $(A^{\oplus_\alpha}2)$ and $(A^{\oplus_\alpha}7)$, the following immediately follows.

$(T^{\oplus_\alpha}1)$ $\mathcal{K} \oplus_\alpha \phi$ is α-regular if and only if $[(\mathbf{R} \uparrow) \wedge (\mathbf{R} \downarrow)] \in \mathrm{Cn}(\{\phi\})$.

The intuition behind $(A^{\oplus_\alpha}6)$ is that consistent supposition in a non-regular context should adhere as much as possible to the revision postulates. $(A^{\oplus_\alpha}7)$ guarantees that, unless forced, supposition always yields a non-regular context. In fact, we can prove the following.

Observation 2 *If $\mathbb{C}$ is an $\mathcal{L}$ context system, such that every $\mathcal{K} \in \mathbb{C}$ is not α-regular, an operator $\odot : \mathbb{C} \times \sigma_P \longrightarrow \mathbb{C}$ is a revision operator if and only if it is an α-supposition operator.*

Example 2. The condition requiring all contexts in $\mathbb{C}$ to be not α-regular is satisfiable by placing restrictions on the underlying language $\mathcal{L}$. Following are examples of such restrictions.

1. $\mathcal{L} = \{\phi|\phi$ does not contain the symbol $\mathbf{B}\}$: In this case, we restrict reasoning to objective sentences only. Neither $(\mathbf{R} \uparrow)$ nor $(\mathbf{R} \downarrow)$ can be in any context.
2. $\mathcal{L} = \{\phi|\phi$ does not contain the symbol $\overrightarrow{\mathbf{B}}\}$: Here, the function symbol $\overrightarrow{\mathbf{B}}$ is not part of the language. Thus, $(\mathbf{R} \uparrow)$ cannot be in any context.
3. $\mathcal{L} = \{\phi|\phi$ does not contain $\mathbf{B}(\alpha, \psi)$ as a sub-term for any term $\psi\}$: Since $\mathbf{B}(\alpha, \psi)$ is not mentionable in the language, then so are $(\mathbf{R} \uparrow)$ and $(\mathbf{R} \downarrow)$.

□

Thus, for an autodoxastic agent α, the above example shows that revision and α-supposition are indistinguishable under at least the following conditions: (i) α cannot reason about belief, (ii) α can reason about belief but not about introspection, and (iii) α can reason about beliefs and introspection but not about its own beliefs.

7   Varieties of Supposition

The postulates of supposition are general enough to admit a number of different supposition operators, each with a different intuitive flavor. Before defining these operators, however, one notion is needed.

Definition 2 *For a context $\mathcal{K}$ and an $\mathcal{L}^P$ term ϕ, the set of α-deregulation subsets of $\mathcal{K}$ with respect to ϕ, is the set*

$$\Delta_\alpha(\mathcal{K}, \phi) = \{\mathcal{K}'|\mathcal{K}' \subseteq \mathcal{K} \text{ and } \mathcal{K}' \cup \{\phi\} \text{ is not } \alpha-\text{regular}\}$$

The set of α-deregulation closures of $\mathcal{K}$ with respect to ϕ is the set

$$\delta_\alpha(\phi, \mathcal{K}) = \{\mathcal{K}' \in \Delta_\alpha(\mathcal{K}, \phi)| \text{ there is no } \mathcal{K}'' \in \Delta_\alpha(\mathcal{K}, \phi) \text{ such that } \mathcal{K}' \subset \mathcal{K}''\}$$

The following observations about deregulation subsets immediately follow.

Observation 3 *Let $\mathcal{K}$ be an $\mathcal{L}$ context, $\phi \in \mathcal{L}^P$, and $\alpha \in \mathcal{L}^A$.*

1. *If $\mathcal{K}' \in \Delta_\alpha(\mathcal{K}, \phi)$, then $\mathcal{K}'$ is not α-regular.*
2. *If $\mathcal{K} \in \Delta_\alpha(\mathcal{K}, \phi)$, then $\delta_\alpha(\phi, \mathcal{K}) = \{\mathcal{K}\}$.*
3. *$\Delta_\alpha(\phi, \mathcal{K}) = \varnothing$ if and only if $[(\mathbf{R}\uparrow) \wedge (\mathbf{R}\downarrow)] \in \mathrm{Cn}(\{\phi\})$.*
4. *If $\phi \in \mathcal{K}$, then $\mathrm{Cn}(\mathcal{K}' \circledast \phi) = \mathrm{Cn}(\mathcal{K}' \cup \{\phi\}) = \mathrm{Cn}(\mathcal{K}')$, for every $\mathcal{K}' \in \delta_\alpha(\phi, \mathcal{K})$.*

Using the notion of deregulation closures, we may now define particular supposition operators. Two of these are especially interesting, we call them conservative post-deregulation and conservative pre-deregulation operators. These operators are conservative in the sense that they preserve as much of the original context as possible.

Definition 3 *Let $\odot_\alpha : \mathbb{C} \times \mathcal{L}^P \longrightarrow \mathbb{C}$.*

1. *$\odot_\alpha$ is a conservative post-deregulation operator if*

$$\mathcal{K} \odot_\alpha \phi = \begin{cases} \mathcal{K} \circledast \phi & \text{if } \delta_\alpha(\phi, \mathcal{K}) = \varnothing \\ \mathcal{K}' \circledast \phi & \text{otherwise, for some } \mathcal{K}' \in \delta_\alpha(\phi, \mathcal{K}) \end{cases}$$

2. *$\odot_\alpha$ is a conservative pre-deregulation operator if*

$$\mathcal{K} \odot_\alpha \phi = \begin{cases} \mathcal{K} \circledast \phi & \text{if } \delta_\alpha(\phi, \mathcal{K}) = \varnothing \\ \mathcal{K}' \circledast \phi & \text{otherwise, for some } \mathcal{K}' \in \delta_\alpha(\phi, \mathcal{K} \circledast \phi) \end{cases}$$

In what follows, a conservative deregulation operator is either a conservative post-deregulation or a conservative pre-deregulation operator.

Theorem 1. *A conservative deregulation operator is a supposition operator.*

A corresponding pair of deregulation operators are more aggressive than the conservative ones in the method they employ for deregulation.

Definition 4 *Let $\odot_\alpha : \mathbb{C} \times \mathcal{L}^P \longrightarrow \mathbb{C}$.*

1. *$\odot_\alpha$ is an aggressive post-deregulation operator if*

$$\mathcal{K} \odot_\alpha \phi = \begin{cases} \mathcal{K} \circledast \phi & \text{if } \delta_\alpha(\phi, \mathcal{K}) = \varnothing \text{ or } \delta_\alpha(\phi, \mathcal{K}) = \{\mathcal{K}\} \\ \mathcal{K}' \circledast \phi & \text{otherwise, for some } \mathcal{K}' \in \delta_\alpha(\phi, \mathcal{K} \circledast \neg[(\mathbf{R}\uparrow) \wedge (\mathbf{R}\downarrow)]) \end{cases}$$

2. *$\odot_\alpha$ is an aggressive pre-deregulation operator if*

$$\mathcal{K} \odot_\alpha \phi = \begin{cases} \mathcal{K} \circledast \phi & \text{if } \delta_\alpha(\phi, \mathcal{K}) = \varnothing \text{ or } \delta_\alpha(\phi, \mathcal{K}) = \{\mathcal{K}\} \\ \mathcal{K}' \circledast \phi & \text{otherwise, for some } \mathcal{K}' \in \delta_\alpha(\phi, (\mathcal{K} \circledast \phi) \circledast \neg[(\mathbf{R}\uparrow) \wedge (\mathbf{R}\downarrow)]) \end{cases}$$

In what follows, an aggressive deregulation operator is either an aggressive post-deregulation or an aggressive pre-deregulation operator.

Theorem 2. *An aggressive deregulation operator is a supposition operator.*

The four supposition operators presented above are deregulation operators varying along two orthogonal dimensions. Conservative deregulation eliminates α-regularity with as minimal perturbation to the context as possible (where minimality is determined by $\circledast$). Thus, by conservatively supposing ϕ, an agent assumes ϕ to be true, but makes no assumption about the completeness or fallibility of its beliefs. As far as such an agent is concerned, the suppositional context admits the possibility of omniscience and infallibility. By aggressive supposition, on the other hand, the agent consciously adopts that it is either ignorant or fallible.

The second factor distinguishing deregulation operators is whether the first revision precedes or follows the deregulation step. The difference is the following: in post-deregulation, the agent supposes only the truth of ϕ, but not that it believes ϕ; pre-deregulation supposes both that ϕ is true and the agent believes it.

8   Conditionals

With a number of belief change operators at our disposal, we now introduce a class of conditional connectives. Let $\mathcal{L}_1$ be the closure of $\mathcal{L}_0$ under the set of binary P-operators $\{\circledast\!\!\rightarrow\} \cup \{\oplus\!\!\rightarrow_\alpha\}_{\alpha \in \mathcal{L}_0^A}$. For every operator $\odot\!\!\rightarrow$ in this set, we add the following two rules to our inference cannon.

$$(\odot\!\!\rightarrow E)\frac{\phi\odot\!\!\rightarrow \psi}{\psi \in \mathrm{Cn}(\mathcal{K} \odot \phi)}$$

$$(\odot\!\!\rightarrow I)\frac{\psi \in \mathrm{Cn}(\mathcal{K} \odot \phi)}{\phi\odot\!\!\rightarrow \psi}$$

$(\odot\!\!\rightarrow E)$ and $(\odot\!\!\rightarrow I)$ are bridge rules [20] constituting the two directions of the Ramsey test, where the underlying operation of belief change is interpreted as $\odot$.

The following is a simple corollary of Observation 2.

Observation 4 *If $\mathbb{C}$ is an $\mathcal{L}$ context system, such that every $\mathcal{K} \in \mathbb{C}$ is not α-regular, then there is an α-supposition operator, $\oplus_\alpha$, such that, for every $\phi, \psi \in \mathcal{L}^P$, $[\phi\circledast\!\!\rightarrow \psi \equiv \phi\oplus\!\!\rightarrow_\alpha \psi] \in \mathrm{Cn}(\mathcal{K})$, for every $\mathcal{K} \in \mathbb{C}$.*

Of course, the particular α-supposition operator referred to is nothing but $\circledast$ itself. Hence, the equivalence of $\circledast\!\!\rightarrow$ and $\oplus\!\!\rightarrow_\alpha$.

Strong connections between $\circledast\!\!\rightarrow$ and $\oplus\!\!\rightarrow_\alpha$ may be established, if we consider only deregulation operators. First of all, it might be tempting to declare that $\circledast\!\!\rightarrow$ and $\oplus\!\!\rightarrow_\alpha$ are equivalent in contexts which are not α-regular. This, however, is not true; our first result is a negative result.

Non-Theorem 1 *Let $\oplus_\alpha$ be a deregulation operator on an $\mathcal{L}$ context system $\mathbb{C}$. For every $\phi, \psi \in \mathcal{L}^P$ and $\mathcal{K} \in \mathbb{C}$, if $\mathcal{K}$ is not α-regular, then $[\phi\circledast\!\!\rightarrow \psi \equiv \phi\oplus\!\!\rightarrow_\alpha \psi] \in \mathrm{Cn}(\mathcal{K})$.*

Counterexample 1. Let $\mathcal{K} = \{(\mathbf{R}\uparrow), p\}$. Clearly, $\mathcal{K}$ is not α-regular. Also, clearly, $(p \supset (\mathbf{R}\downarrow))\circledast\!\!\rightarrow [(\mathbf{R}\uparrow) \wedge (\mathbf{R}\downarrow)] \in \mathrm{Cn}(\mathcal{K})$. Nevertheless, whatever the deregulation operator we assume, $(p \supset (\mathbf{R}\downarrow))\oplus\!\!\rightarrow_\alpha [(\mathbf{R}\uparrow) \wedge (\mathbf{R}\downarrow)] \notin \mathrm{Cn}(\mathcal{K})$, since $\mathcal{K} \oplus_\alpha (p \supset (\mathbf{R}\downarrow))$ is not α-regular, given $(T^{\oplus_\alpha}1)$. $\qquad\square$

On the background of Observation 4, the above non-theorem shows that, once a language is expressive enough to represent regularity, supposition and revision are not generally equivalent in non-regular contexts. We may, however, achieve positive results by strengthening the conditions of non-regularity.

Theorem 3. *Let $\oplus_\alpha$ be a deregulation operator on an $\mathcal{L}$ context system $\mathbb{C}$. For every $\phi, \psi \in \mathcal{L}^P$ and $\mathcal{K} \in \mathbb{C}$, if $\mathcal{K} \cup \{\phi\}$ is not α-regular, then $[(\phi\circledast\!\!\rightarrow \psi) \equiv (\phi\oplus\!\!\rightarrow_\alpha \psi)] \in \mathrm{Cn}(\mathcal{K})$.*

In other words, deregulation and revision with ϕ are equivalent whenever the mere expansion with ϕ will not result in α-regularity. Here, we rely on a strengthening of the conditions of non-regularity. On the other hand, by weakening the relation between $\circledast\!\!\rightarrow$ and $\oplus\!\!\rightarrow_\alpha$ we get a strong positive result for conservative pre-deregulation.

Theorem 4. *Let $\oplus_\alpha$ be a conservative pre-deregulation operator on an $\mathcal{L}$ context system $\mathbb{C}$. For every $\phi, \psi \in \mathcal{L}^P$ and $\mathcal{K} \in \mathbb{C}$, $[(\phi\oplus\!\!\rightarrow_\alpha \psi) \supset (\phi\circledast\!\!\rightarrow \psi)] \in \mathrm{Cn}(\mathcal{K})$.*

A similar pervasive result cannot be derived, however, for conservative post-deregulation.

Non-Theorem 2 *Let $\oplus_\alpha$ be a conservative post-deregulation operator on an $\mathcal{L}$ context system $\mathbb{C}$. For every $\phi, \psi \in \mathcal{L}^P$ and $\mathcal{K} \in \mathbb{C}$, $[(\phi\oplus\!\!\rightarrow_\alpha \psi) \supset (\phi\circledast\!\!\rightarrow \psi)] \in \mathrm{Cn}(\mathcal{K})$.*

Counterexample 2. Let $\mathcal{K} = \{(\mathbf{R}\uparrow), (\mathbf{R}\downarrow), \neg\mathbf{B}(\alpha, p)\}$. From $(A^\circledast 7)$, it follows that $\mathrm{Cn}(\mathcal{K} \circledast p) = \mathrm{Cn}(\{(\mathbf{R}\uparrow), (\mathbf{R}\downarrow), p\})$. On the other hand, $\mathrm{Cn}(\mathcal{K} \oplus_\alpha p) = \mathrm{Cn}(\{p, \neg\mathbf{B}(\alpha, p)\})$. $\qquad\square$

Similar to conservative post-deregulation, no assumption weaker than the non-regularity of $\mathcal{K} \cup \{\phi\}$ may guarantee any implications between $\circledast\!\!\rightarrow$ and aggressive $\oplus\!\!\rightarrow$. Basically, since aggressive supposition involves an initial revision with $\neg[(\mathbf{R}\uparrow) \wedge (\mathbf{R}\downarrow)]$, $\mathrm{Cn}(\mathcal{K} \oplus_\alpha \phi) \subseteq \mathrm{Cn}(\mathcal{K} \circledast \phi)$ is, in general, not true. On the other hand, since $\mathcal{K} \circledast \phi$ may be α-regular while $\mathcal{K} \oplus_\alpha \phi$ is usually not, $\mathrm{Cn}(\mathcal{K} \circledast \phi) \subseteq \mathrm{Cn}(\mathcal{K} \oplus_\alpha \phi)$ is not guaranteed either.

9   Conclusions

Despite the apparent tension between the Ramsey test and Thomason conditionals, a suppositional interpretation of the Ramsey test resolves the conflict. We have presented two classes of conditionals, the epistemic conditionals based on belief revision, and the suppositional conditionals based on supposition. Four

supposition operators where defined in terms of revision. These operators give rise to different classes of suppositional conditionals, each with its own features. Relations between these conditionals and the corresponding epistemic conditional, we believe, conform with our intuitions. In particular, we have shown that, in the absence of regular contexts, revision and supposition are identical. This explains the elusiveness of the distinction between the epistemic and the suppositional readings of the Ramsey test, which is only revealed in autodoxastic contexts where non-regular contexts are admitted.

The work presented here is in no way complete: formal semantics for $Log_A\mathbf{B}$ expanded with conditionals is yet to be elucidated; the relation between revision, supposition, and update [21] awaits investigation; and a non-monotonic treatment of autodoxastic logic seems to be the next natural step.

References

1. Stalnaker, R.: A theory of conditionals. In Rescher, N., ed.: Studies in Logical Theory. Number 2 in American Philosophical Quarterly, Monograph Series. Basil Blackwell Publishers, Oxford (1968) 98–112
2. Alchourron, C.E., Gärdenfors, P., Makinson, D.: On the logic of theory change: Partial meet contraction and revision functions. The Journal of Symbolic Logic **50**(2) (1985) 510–530
3. Nebel, B.: A knowledge level analysis of belief revision. In Brachman, R., Levesque, H., Reiter, R., eds.: Principles of Knowledge Representation and Reasoning: Proceedings of the First International Conference (KR'89), Toronto, ON, Morgan Kaufmann (1989) 301–311
4. Williams, M.A.: Transmutations of knowledge systems. In Doyle, J., Sandewall, E., Torasso, P., eds.: Principles of Knowledge Representation and Reasoning: Proceedings of the Fourth International Conference (KR'94), San Francisco, CA, Morgan Kaufmann (1994) 619–629
5. Gärdenfors, P.: Belief revisions and the Ramsey test for conditionals. The Philosophical Review **95**(1) (1986) 81–93
6. Ismail, H.O.: A reason maintenance perspective on relevant Ramsey conditionals. Logic Journal of the IGPL **18**(4) (2010) 508–529
7. van Fraassen, B.: Review of Brian Ellis, Relational belief systems. Canadian Journal of Philosophy **10**(3) (1980) 497–511
8. Chalmers, D., Hájek, A.: Ramsey + Moore = God. Analysis **67** (2007) 170–172
9. Barnett, D.: Ramsey + Moore ≠ God. Analysis **68** (2008) 168–174
10. Willer, M.: New surprises for the Ramsey test. Synthese **176** (2010) 291–309
11. Hannes, L.: God − Moore = Ramsey (a reply to Chalmers and Hájek). Topoi **30**(1) (2011) 47–51
12. Moore, R.C.: Semantic considerations on nonmonotonic logic. Artificial Intelligence **25** (1985) 75–94
13. Perlis, D.: Languages with self-refernce I: Foundations. Artificial Intelligence **25** (1985) 301–322
14. Ismail, H.O.: $Log_A\mathbf{B}$: An algebraic logic of belief. In Beierle, C., Kern-Isberner, G., eds.: Proceedings of the Workshop on Relational Approaches to Knowledge Representation and Learning: Workshop of the 32nd German Conference on Artificial Intelligence. (2009) 2–18

15. Ismail, H.O.: $Log_A\mathbf{B}$: A first-order, non-paradoxical, algebraic logic of belief. Logic Journal of the IGPL (In Press)
16. Burris, S., Sankappanavar, H.P.: A Course in Universal Algebra. Springer-Verlag (1982)
17. Hintikka, J.: Knowledge and Belief: An Introduction to the Logic of the Two Notions. Cornell University Press, Ithaca, New York (1962)
18. Fagin, R., Halpern, J., Moses, Y., Vardi, M.: Reasoning about Knowledge. The MIT Press, Cambridge, Massachusetts (1995)
19. Kaplan, D.: Quantifying in. Synthese **19** (1968) 178–214
20. Giunchiglia, F., Serafini, L.: Multilanguage hierarchial logics (or: How we can do without modal logics). Artficial Intelligence **65** (1994) 29–70
21. Katsuno, H., Mendelzon, A.: On the difference between updating a knowledge base and revising it. In: Principles of Knowledge Representation and Reasoning: Proceedings of the second International Conference (KR'91). (1991) 387–394

A   Proofs

A.1   Proof of Theorem 1

Let $\odot_\alpha$ be a conservative post-deregulation operator. Since $\mathcal{K} \odot_\alpha \phi$ is basically the result of revising some context with ϕ, then $\odot_\alpha$ satisfies $(A^{\oplus_\alpha}1)$ through $(A^{\oplus_\alpha}4)$.

$(\mathbf{A}^{\oplus_\alpha}\mathbf{5})$. If $\mathcal{K} \cup \{\phi\}$ is not α-regular, then by the second clause of Observation 3, $\delta_\alpha(\mathcal{K}, \phi) = \mathcal{K}$. Hence, $\mathcal{K} \odot_\alpha \phi = \mathcal{K} \circledast \phi$. Thus, by $(A^{\circledast}4)$, $\odot_\alpha$ satisfies $(A^{\oplus_\alpha}5)$.

$(\mathbf{A}^{\oplus_\alpha}\mathbf{6})$. If $\delta_\alpha(\mathcal{K}, \phi) = \varnothing$, then $\mathcal{K} \cup \{\phi\}$ is α-regular, and $(A^{\oplus_\alpha}6)$ is trivially true of $\odot_\alpha$. Otherwise, let $\mathcal{K}' \in \delta_\alpha(\mathcal{K}, \phi)$. If $\mathcal{K} \cup \{\phi\}$ is not α-regular, then $\mathcal{K}' \in \delta_\alpha(\mathcal{K}, \phi)$, and, by the second clause of Observation 3, $\mathcal{K}' = \mathcal{K}$. Hence, $\mathcal{K} \odot_\alpha \phi = \mathcal{K} \circledast \phi$ and, by $(A^{\circledast}4)$, $\mathrm{Cn}(\mathcal{K} \cup \{\phi\} \subseteq \mathrm{Cn}(\mathcal{K} \odot_\alpha \phi))$.

$(\mathbf{A}^{\oplus_\alpha}\mathbf{7})$. If $\delta_\alpha(\mathcal{K}, \phi) = \varnothing$, then $\mathcal{K} \odot_\alpha \phi = \mathcal{K} \circledast \phi$. Moreover, $\Delta_\alpha(\mathcal{K}, \phi) = \varnothing$ and, by the third clause of Observation 3, $[(\mathbf{R}{\uparrow}) \wedge (\mathbf{R}{\downarrow})] \in \mathrm{Cn}(\{\phi\})$. Hence, $(A^{\oplus_\alpha}7)$. On the other hand, if $\delta_\alpha(\mathcal{K}, \phi) \neq \varnothing$, then $\mathcal{K} \odot_\alpha \phi = \mathcal{K}' \circledast \phi$, for some $\mathcal{K}' \in \delta_\alpha(\mathcal{K}, \phi)$. Since $\mathcal{K}' \cup \{\phi\}$ is not α-regular, then, by $(A^{\circledast}4)$ and the monotonicity of Cn, it follows that $\mathcal{K} \odot_\alpha \phi = \mathcal{K}' \circledast \phi$ is not α-regular either. Thus, $(A^{\oplus_\alpha}7)$ trivially follows.

Now, let $\odot_\alpha$ be a conservative post-deregulation operator. Similar to the post-regulation case, $\odot_\alpha$ satisfies $(A^{\oplus_\alpha}1)$ through $(A^{\oplus_\alpha}4)$.

$(\mathbf{A}^{\oplus_\alpha}\mathbf{5})$. If $\delta_\alpha(\mathcal{K}, \phi) = \varnothing$, then $\mathcal{K} \odot_\alpha \phi = \mathcal{K} \circledast \phi$, which, by $(A^{\circledast}4)$, satisfies $(A^{\oplus_\alpha}5)$. Otherwise, $\mathcal{K} \odot_\alpha \phi = \mathcal{K}' \circledast \phi$, for some $\mathcal{K}' \in \delta_\alpha(\mathcal{K} \circledast \phi, \phi)$. Since $\phi \in \mathcal{K} \circledast \phi$ (by $(A^{\circledast}2)$), then, from the fourth clause of Observation 3, it follows that $\mathrm{Cn}(\mathcal{K} \odot_\alpha \phi) = \mathrm{Cn}(\mathcal{K}' \circledast \phi) = \mathrm{Cn}(\mathcal{K}' \cup \{\phi\})$. But by $(A^{\circledast}2)$, $(A^{\circledast}4)$, and monotonicity of Cn, $\mathrm{Cn}(\mathcal{K}' \cup \{\phi\}) \subseteq \mathrm{Cn}((\mathcal{K} \circledast \phi) \cup \{\phi\}) = \mathrm{Cn}(\mathcal{K} \circledast \phi) \subseteq \mathrm{Cn}(\mathcal{K} \cup \{\phi\})$. Hence, $(A^{\oplus_\alpha}5)$.[7]

---

[7] In this proof, we did not make use of the assumption that $\mathcal{K} \cup \{\alpha\}$ is not α-regular. An alternative proof goes along the lines of the corresponding proof for conservative post-deregulation.

($\mathbf{A}^{\oplus_\alpha}\mathbf{6}$). If $\delta_\alpha(\mathcal{K}, \phi) = \varnothing$, then $\mathcal{K} \cup \{\phi\}$ is α-regular, and ($A^{\oplus_\alpha}6$) is trivially true of $\odot_\alpha$. Otherwise, let $\mathcal{K}' \in \delta_\alpha(\mathcal{K} \circledast \phi, \phi)$. If $\mathcal{K} \cup \{\phi\}$ is not α-regular, then, by ($A^{\circledast}4$) and monotonicity of Cn, so is $\mathcal{K} \circledast \phi$. But by ($A^{\circledast}2$), $\mathcal{K} \circledast \phi = (\mathcal{K} \circledast \phi) \cup \{\phi\}$. Thus, $\mathcal{K} \circledast \phi \in \delta_\alpha(\mathcal{K} \circledast \phi, \phi)$, and, by the second clause of Observation 3, $\mathcal{K}' = \mathcal{K} \circledast \phi$. Hence, $\mathcal{K} \odot_\alpha \phi = (\mathcal{K} \circledast \phi) \circledast \phi$ and, by ($A^{\circledast}2$) through ($A^{\circledast}5$), $\text{Cn}(\mathcal{K} \odot_\alpha \phi) = \text{Cn}((\mathcal{K} \circledast \phi) \circledast \phi) = \text{Cn}(\mathcal{K} \circledast \phi)$. Thus, by ($A^{\circledast}5$), ($A^{\oplus_\alpha}6$) follows.

($\mathbf{A}^{\oplus_\alpha}\mathbf{7}$). The proof is similar, *mutatis mutandis*, to the proof in the case of a conservative post-deregulation operator.   □

A.2   Proof of Theorem 2

For limitations of space, we only prove the case of aggressive post-deregulation operators. ($A^{\oplus_\alpha}1$) through ($A^{\oplus_\alpha}6$) are obviously satisfied.

To prove ($A^{\oplus_\alpha}7$), consider three cases. First, if $\delta_\alpha(\mathcal{K}, \phi) = \varnothing$, then, by the third clause of Observation 3, $[(\mathbf{R}{\uparrow}) \wedge (\mathbf{R}{\downarrow})] \in \text{Cn}(\phi)$, which implies ($A^{\oplus_\alpha}7$). Second, if $\delta_\alpha(\mathcal{K}, \phi) = \{\mathcal{K}\}$, then $\mathcal{K} \cup \{\phi\}$ is not α-regular. By the monotonicity of Cn, $\mathcal{K} \odot_\alpha \phi = \mathcal{K} \circledast \phi$ is also not α-regular. Thus, ($A^{\oplus_\alpha}7$) is trivially true. Third, $\mathcal{K} \odot_\alpha \phi = \mathcal{K}' \circledast \phi$ for some $\mathcal{K}' \in \delta_\alpha(\phi, \mathcal{K} \circledast \neg[(\mathbf{R}{\uparrow}) \wedge (\mathbf{R}{\downarrow})])$. But, by definition of δ, $\mathcal{K}' \cup \{\phi\}$ is not α-regular. Thus, by monotonicity of Cn, so is $\mathcal{K}' \circledast \phi$, from which ($A^{\oplus_\alpha}7$) trivially follows.

A.3   Proof of Theorem 3

The proof is straightforward for the two cases of aggressive deregulation operators, since the α-non-regularity of $\mathcal{K} \cup \{\phi\}$ implies that $\delta_\alpha(\mathcal{K}, \phi) = \{\mathcal{K}\}$, which in turn implies that $\mathcal{K} \oplus_\alpha \phi = \mathcal{K} \circledast \phi$.

The case of a conservative post-deregulation operator is also straightforward: If $\mathcal{K} \cup \{\alpha\}$ is not α-regular, then $\delta_\alpha(\mathcal{K}, \phi) = \{\mathcal{K}\}$, which implies that $\mathcal{K} \oplus_\alpha \phi = \mathcal{K} \circledast \phi$. For a conservative pre-deregulation operator, note that $\mathcal{K} \cup \{\alpha\}$ being not α-regular implies, by monotonicity of Cn, that $\mathcal{K} \circledast \phi$ is also not α-regular. But, given ($A^{\circledast}2$), $(\mathcal{K} \circledast \phi) \cup \{\alpha\} = \mathcal{K} \circledast \phi$. Thus, $\delta_\alpha(\phi, \mathcal{K} \circledast \phi) = \{\mathcal{K} \circledast \phi\}$. Hence, $\mathcal{K} \oplus_\alpha \phi = (\mathcal{K} \circledast \phi) \circledast \phi = \mathcal{K} \circledast \phi$, by ($A^{\circledast}2$) through ($A^{\circledast}5$).

In all cases, given ($\circledast{\rightarrow} E$), ($\circledast{\rightarrow} I$), ($\oplus{\rightarrow}_\alpha E$), and ($\oplus{\rightarrow}_\alpha I$), it follows that $[(\phi\circledast{\rightarrow} \psi) \equiv (\phi\oplus{\rightarrow}_\alpha \psi)] \in \text{Cn}(\mathcal{K})$.

A.4   Proof of Theorem 4

If $\delta_\alpha(\mathcal{K}, \phi) = \varnothing$, then $\mathcal{K} \oplus_\alpha \phi = \mathcal{K} \circledast \phi$. Otherwise, by the third clause of Observation 3, $[(\mathbf{R}{\uparrow}) \wedge (\mathbf{R}{\downarrow})] \notin \text{Cn}(\{\phi\})$. Thus, there is a $\mathcal{K}' \in \delta_\alpha(\phi, \mathcal{K} \circledast \phi)$ such that $\mathcal{K} \oplus_\alpha \phi = \mathcal{K}' \circledast \phi$. By the fourth clause of Observation 3, $\text{Cn}(\mathcal{K} \oplus_\alpha \phi) = \text{Cn}(\mathcal{K}' \circledast \phi) = \text{Cn}(\mathcal{K}') \subseteq \text{Cn}(\mathcal{K} \circledast \phi)$. Thus, given ($\circledast{\rightarrow} E$), ($\circledast{\rightarrow} I$), ($\oplus{\rightarrow}_\alpha E$), and ($\oplus{\rightarrow}_\alpha I$), it follows that $[(\phi\oplus{\rightarrow}_\alpha \psi) \equiv (\phi\circledast{\rightarrow} \psi)] \in \text{Cn}(\mathcal{K})$.