From Onions to Broccoli: Generalizing Lewis's counterfactual logic[*]

Patrick Girard
Stanford University

May 18, 2005

Abstract

We present a generalization of Segerberg's onion semantics for belief revision, in which the linearity of the spheres need not occur. The resulting logic is called broccoli logic. We provide a minimal relational logic, introducing a new neighborhood semantics operator. We then show that broccoli logic is a well-known conditional logic, the Burgess-Veltman minimal conditional logic.

Belief revision is the study of theory change in which a set of formulas is ascribed to an agent as a belief set revisable in the face of new information. A dominant paradigm in belief revision is the so-called *AGM* paradigm, which describes a functional notion of revision (cf. [1]). A natural semantics in terms of sphere systems (cf. [6]) was given by Grove in [5] and a logical axiomatization was extensively studied by Segerberg (cf. [9, 10, 11, 12] and the forthcoming [13]). The resulting logic is called "dynamic doxastic logic" (*DDL*). A generalization of the *AGM* approach in which revision is taken to be relational rather than functional was first studied by Lindström and Rabinowicz (cf. [7]). Their

motivation was to formalize cases in which an agent may obtain incomparable belief sets after revision with new information. In this paper, we will pursue this generalization and propose a relational belief revision logic. We call the resulting logic "broccoli logic" (*BL*) and the type of revision it depicts "broccoli revision". As it turns out, and this will be the main result of this paper, *BL* already exists, in the guise of what we call "minimal conditional logic", studied by Burgess and Veltman (cf. Burgess [3] and Veltman [14]).

In section 1, we outline onion semantics and the intended generalization to *BL*. In section 2, we give a minimal relational logic (*MRL*) with its complete proof system. The semantics of this logic is in effect a neighborhood semantics (cf. [4]), but we will interpret it terms of revision instead. In section 3, we will propose ways of expending *MRL* to get a complete proof system for *BL* and we will point at a major difficulty in this task, namely to provide an appropriate arrow-condition for generalized selection functions. Finally, section 4 will show how the quest for a generalized selection function, with the promised difficulties inherent in the project, is avoidable by showing that *BL* is equivalent to *MCL*.

1 Onion and Broccoli Logics

This section presents the onion and broccoli semantics. We give definitions of onion and broccoli models and provide the semantics for the broccoli modal operators.

1.1 Onions

An onion is a linearly ordered sphere system that satisfy the limit condition. More precisely,

Definition 1.1. Let U be a nonempty set. An *onion* $O \in \mathcal{P}(U)$ is a linearly ordered set of subsets of U satisfying the following condition (the limit condition): for all $X \subseteq U$:

$$\bigcup O \cap X \neq \emptyset \Rightarrow \exists Z \in O \text{ s.t. } \forall Y \in O(Y \cap X \neq \emptyset \text{ iff } Z \subseteq Y)$$

The limit condition states that every set intersecting an onion intersects a smallest element. Let $O \bullet X = \{Y \in O : Y \cap X \neq \emptyset\}$. We use the notation $Z\mu(O \bullet X)$ to express that Z is a minimal element of the onion O intersecting X, i.e., for all

$Y \in O \bullet X, Z \subseteq Y$. The limit condition can succinctly be written as:

$$\bigcup O \cap X \neq \emptyset \Rightarrow \exists Z \mu(O \bullet X).$$

1.2 Broccoli semantics

We want to pursue a generalization of onion logic by dropping the requirement of linearity and generalizing the limit condition.

Definition 1.2. Let U be a nonempty set. A *broccoli flower* $\mathcal{B} \in \mathcal{P}(U)$ is a set of subsets satisfying a generalized limit condition.

There are two ways to specify the generalized limit condition of definition 1.2. Let $B|X = \{Y \cap X : Y \in \mathcal{B}\}$. For all $X \subseteq U$, if $\bigcup \mathcal{B} \cap X \neq \emptyset$, either:

1. $\exists S \subseteq \mathcal{B}$ s.t. $\forall Y \in \mathcal{B}(Y \cap X \neq \emptyset \Rightarrow \exists Z \in S\,(Z\mu(\mathcal{B} \bullet X) \wedge Z \subseteq Y))$, or
2. $\exists S \subseteq \mathcal{B}$ s.t. $\forall Y \in \mathcal{B}(Y \cap X \neq \emptyset \Rightarrow \exists Z \in S\,((Z \cap X)\mu((\mathcal{B}|X) \bullet X) \wedge Z \subseteq Y))$.

Intuitively, a generalized limit condition states that every set intersecting a broccoli flower intersects every members of a set S of smallest elements of the flower. In the first case, the members of S are minimal sets of the broccoli that have a nonempty intersection with X. In the second case, the members of S have minimal intersection with X.

With a generalized limit condition in hand, it is meaningful to define counterfactual modalities. Two natural candidates for BL (with their respective dual) come to mind. Let '$\rightarrow$' stand for the material conditional. The first modality says that $\varphi \rightarrow \psi$ holds throughout every minimal φ-sphere; the second says that $\varphi \rightarrow \psi$ is consistent with every minimal φ-sphere. We will follow Chellas [4] and write these two counterfactuals as the unary modalities $[\varphi]\psi$ and $[\varphi\rangle\psi$.

Definition 1.3. $\mathfrak{M} = (U, \{\mathcal{B}_u\}_{u \in U}, V)$ is a *broccoli model* if U is a set of worlds, $\{\mathcal{B}_u\}_{u \in U}$ is a family of broccoli flowers for each world $u \in U$ satisfying either generalized limit condition, and V is a valuation assigning sets of worlds to propositions.

In what follows, we suppress the index u whenever it is clear from context.

Definition 1.4. We say that φ is true at world u in a broccoli model $\mathfrak{M}$, written $\mathfrak{M}, u \vDash \varphi$ iff (taking standard truth definition for the propositional and the Boolean cases):

Figure 1: Broccoli semantics of the counterfactual operators $[\varphi]\psi$ and $[\varphi\rangle\psi$, with $|\varphi| = P$ and $|\psi| = Q$.

1. $\mathfrak{M}, u \vDash^{\mathfrak{M}} [\varphi]\psi$ iff $\forall Z\mu(\mathcal{B} \bullet |\varphi|)(Z \cap |\varphi| \subseteq |\psi|)$, and
2. $\mathfrak{M}, u \vDash^{\mathfrak{M}} [\varphi\rangle\psi$ iff $\forall Z\mu(\mathcal{B} \bullet |\varphi|)(Z \cap |\varphi| \cap |\psi|) \neq \emptyset$.

Here, as usual, $|\varphi| = \{u : \mathfrak{M}, u \vDash \varphi\}$. We call $|\varphi|$ the *associated proposition* to φ.

These two modalities are meaningful with either generalized limit condition proposed above. Figure 1 illustrates the semantics of both operators.

2 Minimal relational logic

Our first goal is to get a logic that captures a notion of belief revision in which revision is relational rather than functional. That is, we want to allow for incomparable revisions with respect to a belief set. With that purpose in mind, we need a language that can express notions like "all sets obtained under revision by φ are ψ-sets" and "ψ is consistent with all sets obtained under revision by φ". In counterfactual terminology, the same claims would be read as "all minimal φ-sets are ψ-sets" and "all minimal φ-sets intersect ψ-sets". In this section, we introduce a minimal relational logic that captures the core of these ideas. Section 2.1 introduces the language and the semantics of this minimal logic. We will use intuitive interpretations of the semantics in terms of revision, but this is only to keep the motivation of the paper prominent. We give the axiomatization of the minimal logic in section 2.2 and prove it to be complete in section 2.3.

2.1 Language and Semantics

We use a standard propositional language whose primitive Boolean connectives are negation ¬ and disjunction ∨, augmented with two modalities $[\varphi]\psi$ and $[\varphi\rangle\psi$.

Definition 2.1. Given a finite set of propositional variables P, a *minimal relational model* is a triple (U, R, V), where:

- U is a nonempty set, the universe;
- $R = \{R_{|\varphi|} : \varphi \text{ is a formula}, R_{|\varphi|} \subseteq U \times \mathcal{P}(U)\}$; and
- $V : P \longrightarrow \mathcal{P}(U)$.

Definition 2.2. Let $\mathfrak{M}$ be a model and let $w \in U$. The truth-definition for atomic propositions, negations and disjunction is standard. We say that the formula φ is true at point u in a minimal relational model $\mathfrak{M}$, written $\mathfrak{M}, u \vDash \varphi$ if :

$$\mathfrak{M}, u \vDash [\varphi]\psi \quad \text{iff} \quad \forall X((u, X) \in R_{|\varphi|} \Rightarrow \forall v \in X, \mathfrak{M}, v \vDash \psi)$$
$$\mathfrak{M}, u \vDash [\varphi\rangle\psi \quad \text{iff} \quad \forall X((u, X) \in R_{|\varphi|} \Rightarrow \exists v \in X, \mathfrak{M}, v \vDash \psi))$$

The semantics of the modalities $[\varphi]$ and $[\varphi\rangle$ contains two levels of quantification and should be read in two stages: 1) the left bracket picks out a set of φ-subsets of the universe and 2) the right bracket evaluates where ψ is true in these subsets. Notice that the semantics given by minimal relational models is a neighborhood semantics (cf. [4]). Indeed, the relation R is a relation between worlds and subsets of the universe. The modality $[\varphi]$ is the usual neighborhood universal modality, but indexed with associated propositions $|\varphi|$. It comes with its dual modality $\langle\varphi\rangle$ with the obvious semantics. The interesting addition of our language is the modality $[\varphi\rangle$, which expresses that every set $R_{|\varphi|}$-related to u satisfies ψ in at least one point. In neighborhood terminology, this modality expresses that every φ-neighborhood contains at least one ψ-point. This latter modality also come with its natural dual $\langle\varphi]$. In the remainder of this paper, we shall no longer appeal to neighborhood semantics. We will instead provide an interpretation in terms of revision, but the reader who prefers to think in terms of neighborhood semantics is urged to do so and to see in what respect it is a generalization of this logic.

Figure 2 presents a simple minimal relational model, in which the world u is $R_{|\varphi|}$-related (illustrated with arrows) to the sets of worlds X and Y. If we let $|\psi| = Q$, then ψ is true at every world of X and Y. Hence, according to the minimal semantics of definition 2.2, $[\varphi]\psi$ is true at u. This is


Figure 2: Minimal relational model

enough to illustrate the semantics of our minimal relational logic, but to give a motivation for pursuing this semantics, we illustrate its rôle in *BL*. Figure 3 depicts a broccoli flower consisting of two sets (doted lines) of which X and Y are

Figure 3: Broccoli flower

Figure 4: Intended semantics of the Broccoli revision operators $[\varphi]\psi$ and $[\varphi\rangle\psi$, with $|\varphi| = P$ and $|\psi| = Q$.

subsets. Irrespectively of the generalized limit assumption ultimately adopted, assume that these two sets are minimal φ-sets (or that the sets X and Y have minimal $|\varphi|$-intersections). Our goal is to add restrictions on the relation $R_{|\varphi|}$ in order to get the sets X and Y as two minimal sets returned under revision by φ. This is illustrated in picture 4, where $|\varphi| = P$. In the picture on the left-hand-side, $[\varphi]\psi$ is true at world u, since every set obtained under revision by φ is a ψ-sets. Similarly, $[\varphi\rangle\psi$ is true at u in the right-hand-side picture, since ψ is consistent with every revision by φ.

We see the motivation of the minimal relational logic of the present section. In a full-blown *BL*, either additional restrictions on the relation *R* or so-called 'generalized selection functions' will play the rôle of selecting minimal revised sets. Once these sets are selected, the minimal relational logic of the present section will provide the logic to evaluate what holds in these sets. We will discuss selection function in section 3 below. For the remainder of this section, we will present the logic of minimal relational logic and prove completeness. Our goal is

6

to get a firm grasp of the core of future expansion to *BL*.

2.2 Proof system

The following set of axioms and rules is complete with respect to onion selection models:

Axioms:

1. Classical tautologies
2. $\langle\varphi\rangle\psi \equiv \neg[\varphi]\neg\psi$
3. $\langle\varphi]\psi \equiv \neg\langle\varphi\rangle\neg\psi$
4. $[\varphi](\psi \to \theta) \to ([\varphi]\psi \to [\varphi]\theta)$
5. $\langle\varphi]\psi \to \langle\varphi](\psi \vee \theta)$
6. $[\varphi]\psi \wedge \langle\varphi]\theta \to \langle\varphi](\psi \wedge \theta)$
7. $\neg\langle\varphi]\top \to [\varphi]\psi$

Rules:

1. Modus Ponens.
2. Necessitation for $[\varphi]$ and $[\varphi\rangle$.
3. If φ and φ' are formulas differing only in φ having an occurrence of θ in one place where φ' has an occurrence of θ', and if $\theta \equiv \theta'$ is a theorem, then $\varphi \equiv \varphi'$ is also a theorem.

Rule 3, *substitution of equivalents*, is applied indiscriminately inside or outside the modal operators. We count the presence of 'φ' inside $[\varphi]$ and $[\varphi\rangle$ as occurrences of φ. For example, if $\psi \equiv \theta$, then both $[\varphi]\psi \equiv [\varphi]\theta$ and $[\psi]\alpha \equiv [\theta]\alpha$ are instances of rule 3.

Axioms 2 and 3 provide the dual modalities of $[\varphi]$ and $[\varphi\rangle$ respectively. Axiom 4 is a typical K axiom for the modality $[\varphi]$ and yields modus ponens under the scope of $[\varphi]$.[1] Axioms 5 is a monotonicity axiom for the modality $\langle\varphi\rangle$. Intuitively, if ψ is consistent with some revision by φ, then anything weaker than ψ is also consistent with some revision by φ. Finally, axiom 6 provides a minimal

---

[1]There is no corresponding K axiom for the $[\varphi\rangle$. Consider a model M such that the set $X \subseteq U$ is the only subset of U that is φ-related to the world $u \in U$, i.e, $R_{|\varphi|} = \{(u, X)\}$. Suppose that both $|\psi| \cap X \neq \emptyset$ and $|\neg\psi| \cap X \neq \emptyset$, but that $|\theta| \cap X = \emptyset$. Then $u \vDash [\varphi\rangle(\psi \to \theta)$ (since $|\neg\psi| \cap X \neq \emptyset$) and $u \vDash [\varphi\rangle\psi$, but $u \nvDash [\varphi\rangle\theta$. Hence, $[\varphi\rangle\psi(\psi \to \theta) \to ([\varphi\rangle\psi \to [\varphi\rangle\theta)$ is not valid.

interaction between the modalities: If ψ is consistent with every revision by φ and there is a revision by φ such that θ is consistent, then there is a revision by φ such that both ψ and θ are consistent. Finally, axiom 7 says that if there is no revision by φ, then every $[\varphi]$ formula holds vacuously. This is akin to saying that every necessary formula holds at en end-point in a Kripke model.

Now, Suppose that $\langle\varphi]\top \in u$ for some $u \in U.^2$ Then, for every $\psi \in u$ such that $[\varphi\rangle\psi \in u$, axiom 6 gives that $\langle\varphi](\psi \wedge \top) \in u$. By monotonicity of $\langle\varphi]$ (axiom 5), $\langle\varphi]\psi \in u$. Hence, if there is a revision by φ and if ψ is consistent with every revision by φ, then there is a least one revision by φ that witnesses the consistency of ψ. This is desirable for a belief revision logic.

2.3 Completeness

Soundness is a routine matter. We show the soundness of axiom (6) and leave the others to the reader. Assume that $\mathfrak{M}, u \vDash [\varphi]\psi \wedge \langle\varphi]\theta$. Then $\mathfrak{M}, u \vDash \langle\varphi]\theta$ i.e., $\exists X((u, X) \in R_{|\varphi|} \wedge \forall v \in X, \mathfrak{M}, v \vDash \theta)$. But $\mathfrak{M}, u \vDash [\varphi]\psi$ implies that $\forall v \in X, \mathfrak{M}, v \vDash \psi$. Hence, $\forall v \in X, \mathfrak{M}, v \vDash \psi \wedge \theta$. Therefore, $\exists X((u, X) \in R_{|\varphi|} \wedge \forall v \in X, \mathfrak{M}, v \vDash \psi \wedge \theta$, i.e., $\mathfrak{M}, u \vDash \langle\varphi](\psi \wedge \theta)$.

For the completeness part, let $U^{\mathcal{L}}$ consists of all maximal $\mathcal{L}$-consistent sets of formulas. For each formula φ, we define an accessibility relation $R^{\mathcal{L}}_{|\varphi|}$ between worlds and subsets of worlds of $U^{\mathcal{L}}$. For all world $u \in U^{\mathcal{L}}$, if $\langle\varphi]\top \notin u$, then we put $R^{\mathcal{L}}_{|\varphi|} = \emptyset$. Otherwise, for every subsets $X \subseteq U^{\mathcal{L}}$ and formulas φ and ψ, we say that the ordered pair $(u, X) \in R^{\mathcal{L}}_{|\varphi|}$ iff X satisfies the two following conditions:

1. for all $x \in X$, if $[\varphi]\psi \in u$, then $\psi \in x$; and
2. for every $[\varphi\rangle\psi \in u$, X contains at least one world v with $\psi \in v$.

Definition 2.3. Let $p \in P$ be a proposition. Let $V^{\mathcal{L}}(p) = |p|$ and let $R^{\mathcal{L}} = \{R^{\mathcal{L}}_{|\varphi|} : \varphi \text{ is a formula}\}$, then the model $\mathfrak{M}^{\mathcal{L}} = (U^{\mathcal{L}}, R^{\mathcal{L}}, V^{\mathcal{L}})$ is the *canonical minimal relational model*.

The completeness of the proof system in section 2.2 follows from a standard truth-lemma:

Lemma 2.4. *For all $u \in U^{\mathcal{L}}$ and for all formula θ, $\theta \in u$ iff $\mathfrak{M}, u \vDash \theta$.*

---

[2] We read $\langle\varphi]\top$ as "there is a revision by φ".

We will give the proof of the truth-lemma once we have stated and proved the following crucial lemmas.

Lemma 2.5. *If* $\langle\varphi]\alpha \in u$, *then there exists a subset* $X \subseteq U^{\mathcal{L}}$ *such that* $R^{\mathcal{L}}_{|\varphi|}uX$, *and for every world* $x \in X$, $\alpha \in x$.

Proof. Let $[\varphi\rangle\theta \in u$, and let

$$v^- = \{\beta : [\varphi]\beta \in u\} \cup \{\theta\} \cup \{\alpha\}$$

then v^- is consistent. Suppose that v^- is not consistent, then there exists $\delta_1, ..., \delta_n \in v^-$ such that $\vdash (\bigwedge \delta_i \wedge \alpha) \to \neg\theta$. For every $1 \le i \le n$,

$$
\begin{array}{lll}
\delta_i \in v^- & \Rightarrow & [\varphi]\delta_i \in u & \text{(Definition of } v^-) \\
& \Rightarrow & \bigwedge[\varphi]\delta_i \in u & \text{(Truth definition)} \\
& \Rightarrow & [\varphi]\bigwedge\delta_i \in u & \text{(Axiom 4)} \\
& \Rightarrow & ([\varphi]\bigwedge\delta_i \wedge \langle\varphi]\alpha) \in u & \text{(since } \langle\varphi]\alpha \in u) \\
& \Rightarrow & \langle\varphi](\bigwedge\delta_i \wedge \alpha) \in u & \text{(axiom 6)} \\
& \Rightarrow & \langle\varphi]\neg\theta \in u & \text{(by the monotonicity axiom 5)} \\
& \Rightarrow & \neg[\varphi\rangle\theta \in u & \text{(axiom 3)}
\end{array}
$$

and this is a contradiction, since $[\varphi\rangle\theta \in u$ by assumption. Therefore, v^- is consistent. Let v be a maximal extension of v^-.

For every θ_i such that $[\varphi\rangle\theta_i \in u$, let w_i be obtained from the above construction, and let

$$X = \{w_i : [\varphi\rangle\theta_i \in u, \theta_i \in w_i\}.$$

Then X satisfies conditions (1) and (2) and for every $x \in X$, $\alpha \in x$. □

Corollary 2.6 (Corollary to the proof of lemma 2.5). *If* $[\varphi\rangle\psi \in u$, *then the set* $w = \{\psi\} \cup \{\theta : [\varphi]\theta \in u\}$ *is consistent.*

Lemma 2.7. *If* $\langle\varphi\rangle\psi \in u$, *then there exists a subset* $X \subseteq U^{\mathcal{L}}$ *such that* $R^{\mathcal{L}}_{|\varphi|}uX$, *and there exists a world* $x \in X$ *such that* $\psi \in x$.

Proof. Assume $\langle\varphi\rangle\psi \in u$. Then there is a maximal consistent set v such that $\psi \in v$. The proof that v exists is standard (see [2], Lemma 4.20).

Now, by corollary 2.6, for every formula α_i, if $[\varphi\rangle\alpha_i \in u$, then the set $w_i^- = \{\alpha_i\} \cup \{\theta : [\varphi]\theta\}$ is consistent. By Lindenbaum's lemma, there exists a maximal consistent set w_i extending w_i^- such that $\alpha_i \in w_i$. Let W be the set of the w_i such that $[\varphi\rangle\alpha_i \in u$.

Finally, let $X = \{v\} \cup W$. It is not difficult to check that $R^{\mathcal{L}}_{|\varphi|}uX$, and $\psi \in v$. □

We are now ready for the proof of the truth-lemma.

Proof of Lemma 2.4. Thanks to axioms 5 and 7, if $\langle\varphi\rangle\top \notin u$, then $[\varphi\rangle\psi \in u$ and $[\varphi]\psi \in u$ for all ψ. And since $R^{\mathcal{L}}_{|\varphi|} = \emptyset$ when $\langle\varphi\rangle\top \notin u$, the lemma is trivially satisfied. Thus, we assume for the remainder of the proof that $\langle\varphi\rangle\top \in u$. The proof now proceeds by induction on the complexity of θ. The interesting cases are when $\theta = [\varphi]\psi$ or $\theta = [\varphi\rangle\psi$. The first direction ($\theta \in u \Rightarrow \mathfrak{M}, u \vDash \theta$) follows from the conditions imposed on $R^{\mathcal{L}}_{|\varphi|}$. We prove that $\mathfrak{M}, u \vDash \theta \Rightarrow \theta \in u$.

Suppose $[\varphi]\psi \notin u$. Since u is a maximal consistent set of formulas, $\neg[\varphi]\psi \in u$. By axiom 2, this implies that $\langle\varphi\rangle\neg\psi \in u$. By lemma 2.7, there exists a subset $x \subseteq U^{\mathcal{L}}$ such that $R^{\mathcal{L}}_{|\varphi|}uX$ and a world $x \in X$ such that $\mathfrak{M}, x \vDash \neg\psi$. Hence, by truth-definition $\mathfrak{M}, u \vDash \langle\varphi\rangle\neg\psi$, i.e., $\mathfrak{M}, u \vDash \neg[\varphi]\psi$. Therefore, $\mathfrak{M}, u \nvDash [\varphi]\psi$

Finally, suppose that $[\varphi\rangle\psi \notin u$, then $\neg[\varphi\rangle\psi \in u$. Hence, $\langle\varphi\rangle\neg\psi \in u$ (axiom 3). By lemma 2.5, there exists a subset $X \subseteq U^{\mathcal{L}}$ such that $R^{\mathcal{L}}_{|\varphi|}uX$ and for every world $x \in X$, $\neg\psi \in x$. By inductive hypothesis, for every $x \in X$, $\mathfrak{M}, x \vDash \neg\psi$. Hence, by truth-definition, $\mathfrak{M}, u \nvDash [\varphi\rangle\psi$. $\square$

3 Generalized Selection Functions

Selection functions are a natural level in between onion semantics and the general neighborhood models of the preceding section. A selection function takes a proposition p as an argument and return a segment from the set of closest p-worlds (cf. [6]). In this section, we show what properties a suitably generalized selection function should satisfy to play the same rôle as selection functions in onion semantics, and we point to the difficulties of the generalization. We start with selection functions for the original case of onion models.

3.1 Onions and selection functions

Definition 3.1. A function $f : \mathcal{P}(U) \to \mathcal{P}(U)$ is a *selection function* if it satisfies the following conditions, where $X \subseteq U$:

$$f(x) \subseteq x \qquad\qquad\qquad\qquad\qquad\qquad \text{(INC)}$$
$$x \subseteq y \Rightarrow (f(x) \neq \emptyset \Rightarrow f(y) \neq \emptyset) \qquad \text{(MON)}$$
$$x \subseteq y \Rightarrow (X \cap f(y) \neq \emptyset \Rightarrow f(x) = X \cap f(y)) \quad \text{(ARR)}$$

The third condition is called the *arrow condition* and is the source of the major difficulties in the original development of broccoli semantics.

Let U be a finite set and let F be a selection function on U. Let

$$
\begin{aligned}
S_0 &= F(U) \\
S_{n+1} &= S_n \cup F(U - S_n)
\end{aligned}
$$

Since U is finite, there is a smallest m such that $S_{m+1} = S_m$. We leave to the reader to verify that the set $O_F = \{S_n : n < m\}$ is an onion and that O_F and F agree.[3] Hence, models for onions may given in terms of selection functions.

Definition 3.2. Let U be a set, F a selection function on U and V a valuation on a given set of propositional variables. We say that $\mathfrak{M} = (U, F, V)$ is an *onion selection model*.

The truth-definition for the modality $[\varphi]$ in onion selection model is given by:

$$
\mathfrak{M}, u \vDash [\varphi]\psi \text{ iff } F_u(|\varphi|) \subseteq |\psi|.
$$

The complete logic for onions consists of the axioms (1), (2), and (4) of section 2.2 together with the additional axioms $(I), (M)$ and (A):

$$
\begin{aligned}
&[\varphi]\varphi &&(I) \\
&\langle\varphi\rangle\psi \rightarrow \langle\psi\rangle T &&(M) \\
&\langle\varphi\rangle\psi \rightarrow ([\varphi \wedge \psi]\theta \equiv [\varphi](\psi \rightarrow \theta)) &&(A)
\end{aligned}
$$

Axioms $(I), (M)$ and (A) are obvious analogues of conditions $(INC), (MON)$ and (ARR) of definition 3.1. The total resulting system is almost Lewis's famous conditional logic VC, provided that we add an assumption of centrality (cf. [6, 8]).

3.2 Broccoli and generalized selection functions

Now, consider the issue of generalizing this format in a non-linear broccoli setting.

(INC) and (MON) are easily generalized in BL to the following conditions, for all $X, Y \subseteq U$:

$$
\begin{aligned}
&Y \in F(X) \Rightarrow Y \subseteq X &&(INC^*) \\
&Y \subseteq X \text{ and } \exists Z \in F(Y) \text{ s.t. } Z \neq \emptyset \Rightarrow \exists Z \in F(X) \text{ s.t. } Z \neq \emptyset) &&(MON^*)
\end{aligned}
$$

---

[3]O_F and F agree iff

1. $O_F \cap X \neq \emptyset \Rightarrow FX = X \cap S_k$ for some k.
2. $O_F \cap X = \emptyset \Rightarrow FX = \emptyset$.

with the identical corresponding axioms (I) and (M). On the one hand, if $\neg\langle\varphi\rangle\top \notin u$ for some world $u \in U$ (i.e. if there is no revision by φ) then $[\varphi]\varphi \in u$ by axiom 7. But if there is no revision by φ, then $F(X)$ is empty, and (INC^*) holds vacuously. On the other hand, if there is a revision by φ, then (I) and (INC^*) express the same thing. Similar considerations will convince the reader that (M) and (MON^*) go together.

A difficulty arises when attempting to generalize condition (ARR) in a similar fashion. The choice of the generalization depends on the generalized limit condition adopted. For example, a generalization of (ARR) that seems natural is:

3. $Y \subseteq X \Rightarrow (\exists Z \in F(X) \text{ s.t. } Y \cap Z \neq \emptyset \Rightarrow Y \cap Z \in F(Y))$ (ARR^*)

but (ARR^*) holds only if we insisted on keeping the generalized limit condition 1 of section 1.[4] Now, only one half of (A) can be kept in BL, viz. $\langle\varphi\rangle\psi \rightarrow ([\varphi \wedge \psi]\theta \rightarrow [\varphi](\psi \rightarrow \theta))$. The other half makes a crucial appeal to linearity, as may be seen from the counter-model of picture 5. Furthermore, this counter-model invalidates $\langle\varphi\rangle\psi \rightarrow ([\varphi](\psi \rightarrow \theta) \rightarrow [\varphi \wedge \psi]\theta)$ under both limit conditions. It is an open question to find an appropriate axiom that corresponds to condition (ARR^*), or alternatively, to find an appropriate generalization of (ARR) that yields a generalized selection function for BL. This promises to be a difficult task. But instead of pursuing this enterprize further, we pause and see whether BL may not be obtained from an entirely different approach, viz. by showing that the logic already exists! The fact that it is is the third and final contribution of this paper.

4   Broccoli Logic and Minimal Conditional Logic

Minimal conditional logic (MCL) was studied by Stalnaker, Pollock, Burgess and Veltman to capture the idea that a conditional $\varphi \Rightarrow \psi$ is true if an only if the conjunction $\varphi \wedge \neg\theta$ is less possible than the conjunction $\varphi \wedge \theta$, and no more. Their modeling comes with a reflexive and transitive $\leq$-order for each world x and no spheres need occur. In a sphere system, two worlds lying on the same sphere agree on which worlds are farther away and which are closer. This assumption is dropped in MCL. Hence, if two worlds x and y are equally far away in the underlying order from the real world u and if the world z is farther away then the world y, no conclusions may be drawn as to whether world z is farther from the real world then the world x, or vice versa. Instead of changing the onion picture

---

[4]We shall make such an assumption for the remainder of the paper, unless if stated otherwise.

Figure 5: Counter-model to $\langle\varphi\rangle\psi \to ([\varphi](\psi \to \theta) \to [\varphi \wedge \psi]\theta)$, with $P = |\theta|$, $Q = |\varphi|$ and $R = |\psi|$.

by allowing non-linearly ordered sphere system as we wish to do in *BL*, *MCL* ignores spheres altogether. It has been a difficult task to find completeness for *MCL*, and we refer the reader to Burgess [3] for a detailed proof. This section will show how to avoid similar completeness difficulties with *BL* by showing that it is actually equivalent to *MCL*.

Section 4.1 provides the *MCL* semantics, section 4.2 gives the complete proof system and section 4.3 shows the equivalence of *MCL* and *BL*.

4.1   Minimal Conditional Logic

A *Minimal conditional logic model* is a triple (U, R^3, V), where U and V are as above, and R^3 is a ternary relation on U that respects reflexivity and transitivity (cf. [3]). The relation $Rxyz$ should be read as "according to world x, world y is no farther away than world z". We shall write the more suggestive $y \leq_x z$ instead of $Rxyz$. We let $W_u = \{y : \exists z, y \leq_x z\}$ be the zone of entertainability for world $u \in U$. Intuitively, worlds outside the zone of entertainability for u are worlds so far away that their distance from the real world is not appreciable. The *minimal conditional logic language* contains a set of propositional variables, together with negation $\neg$, disjunction $\vee$ and a counterfactual modality $[\varphi]$ for every formula φ.

Definition 4.1. We say that the formula $[\varphi]\psi$ is true at world u in the model $\mathfrak{M}$, and we write $\mathfrak{M}, u \vDash [\varphi]\psi$, iff:

$$\forall y \in W_u \cap V(\varphi), \exists z \in W_u \cap V(\varphi)[z \leq_u y \& \forall w \in W_u \cap V(\varphi)(w \leq_u z \to w \in V(\psi)]$$

13

Notice that the semantic definition of $[\varphi]\psi$ does not contain a minimality condition. However, if the model is finite and $\mathfrak{M}, u \vDash [\varphi]\psi$, then there is a minimal world $z \in U$ such that $z \in V(\varphi) \cap V(\psi)$. Since we will only use finite models for our equivalence result, we use the minimality formulation in evaluating $[\varphi]\psi$ for the remainder of this paper. The semantic condition reduces to:

$$\forall y \in W_U \cap V(\varphi), \exists z \in W_u \cap V(\varphi)[z \leq_u y \& \forall w <_u z, w \notin V(\varphi) \& z \in V(\psi)].$$

4.2 Proof System

The following set of axioms, with the same set of rules as for minimal relational logic presented in section 2.2, is complete for *MCL* (cf. [3]):

1. Classical tautologies
2. $[\varphi]\varphi$
3. $[\varphi]\psi \wedge [\varphi]\theta \to [\varphi](\psi \wedge \theta)$
4. $[\varphi](\psi \wedge \theta) \to [\varphi]\psi$
5. $[\varphi]\psi \wedge [\varphi]\theta \to [\varphi \wedge \psi]\theta$
6. $[\varphi]\psi \wedge [\theta]\psi \to [\varphi \vee \theta]\psi$

We give some examples of derivable theses.

Example 4.2. $MCL \vdash [\varphi]\psi \wedge [\varphi \wedge \psi]\theta \to [\varphi]\theta$

Proof. Assume 1) $\vdash [\varphi]\psi$ and 2) $\vdash [\varphi \wedge \psi]\theta$. By axiom (2), $\vdash [\varphi \wedge \neg\psi](\varphi \wedge \neg\psi)$ and by axiom (4), $\vdash [\varphi \wedge \neg\psi]\neg\psi$. Hence, by monotonicity in the consequent (axiom (4) again), $\vdash [\varphi \wedge \neg\psi](\neg\psi \vee \theta)$. Now, from assumption 2) and axiom (4), $\vdash [\varphi \wedge \psi](\neg\psi \vee \theta)$. Combining the two latter result, we get that $\vdash [\varphi](\neg\psi \vee \theta)$. But since $\vdash [\varphi]\psi$ by assumption (1), we get that $[\varphi]\theta$, as desired. $\qquad\square$

Example 4.3. $MCL \vdash \langle\varphi\rangle\psi \to \langle\psi\rangle T$

Proof. We prove the contrapositive. Assume that $\vdash [\psi]\bot$. Then both $\vdash [\psi]\neg\psi$ and $\vdash [\psi]\varphi$. Hence, by axiom (5), $\vdash [\psi \wedge \varphi]\neg\psi$. But $\vdash [\neg\psi \wedge \varphi](\neg\psi \wedge \varphi)$ is an instance of axiom (2) and by axiom (4), $\vdash [\neg\psi \wedge \varphi]\neg\psi$. Therefore, $\vdash [\varphi]\neg\psi$. $\qquad\square$

Example 4.4. $MCL \vdash [\varphi \wedge \psi]\theta \to [\varphi](\psi \to \theta)$.

Proof. Assume $\vdash [\varphi \wedge \psi]\theta$. By monotonicity, $\vdash [\varphi \wedge \psi](\neg\psi \vee \theta)$. But $\vdash [\varphi \wedge \neg\psi](\neg\psi \vee \theta)$. Therefore, $\vdash [\varphi](\neg\psi \vee \theta)$, i.e., $\vdash [\varphi](\psi \to \theta)$. $\qquad\square$

As we can see from examples axiom (2) and examples 4.3 and 4.4, conditions (*I*), (*M*) and one direction of *A* of section 3 are derivable in *MCL*. We see at once that *MCL* has the properties we were looking for in *BL* (cf. 3), and we will now show that it can get *all* properties of *BL*. The general reason behind these considerations becomes clear in the next subsection.

4.3 Minimal Conditional Logic is Broccoli Logic

Let $\mathfrak{M} = (U, R, V)$ be a finite minimal conditional logic model. We will transform this model into a broccoli model, by constructing a broccoli at each world of $\mathfrak{M}$, taking the downward closed sets of worlds according to the underlying order. More precisely, let $C_x(y) = \{z \in U : z \leq_x y\}$, then $BROC(x) = \{C_x(y) : y \in W_x\}$ is the induced broccoli at world x. In particular, since $\mathfrak{M}$ is finite, the generalized limit condition of definition 1.2 holds. An induced broccoli model $BROC(\mathfrak{M})$ is then given by:

$$BROC(\mathfrak{M}) = \{BROC(x) : x \in U\}$$

The semantics of $[\varphi]\psi$ in the induced broccoli model is given by the following:

$$BROC(\mathfrak{M}), x \vDash [\varphi]\psi \text{ iff } \forall Z\mu(BROC(x) \bullet |\varphi|)(Z \cap |\varphi| \subseteq |\psi|).$$

The main result of this section now follows from lemma 4.5.

Lemma 4.5. $\mathfrak{M}, x \vDash [\varphi]\psi$ *iff* $BROC(\mathfrak{M}), x \vDash [\varphi]\psi$.

Proof. In the one direction, assume that $\mathfrak{M}, x \vDash [\varphi]\psi$. Let $C_w\mu(BROC(x) \bullet |\varphi|)$, and let $v \in C_w \cap |\varphi|$. By the truth definition for $[\varphi]\psi$, $\exists z \leq_x v$ such that $\mathfrak{M}, z \vDash \varphi \wedge \psi$ and $\forall y <_x z, \mathfrak{M}, y \nvDash \varphi$. But z must be equal to v. Otherwise, $C_z \subset C_v \subseteq C_w$ (the latter inclusion uses the transitivity of $\leq_x$), which implies that C_z would be a proper subset of C_w intersecting $|\varphi|$, contradicting our assumption. Thus, $v \in |\psi|$, which implies that $C_w \cap |\varphi| \subseteq |\psi|$. Therefore, as C_w was chosen arbitrarily, $BROC(\mathfrak{M}), x \vDash [\varphi]\psi$.

In the other direction, assume that $BROC(\mathfrak{M}), x \vDash [\varphi]\psi$ and suppose that $\mathfrak{M}, y \vDash \varphi$ for some $y \in U$. Then $C_y \cap |\varphi| \neq \emptyset$. Hence, $\exists C_w \subseteq C_y$ such that $C_w\mu(BROC(x) \bullet |\varphi|)$ (by the limit condition!) and $C_w \cap |\varphi| \subseteq |\psi|$. But since $C_w \subseteq C_y, w \leq_x y$. Assume that w is not a minimal world satisfying $\varphi \wedge \psi$ with respect to $\leq_x$, then $\exists w' <_x w$ such that $\mathfrak{M}, w' \vDash \varphi \wedge \psi$. This implies that $C'_w \subset C_w$ and $C'_w \cap |\varphi| \cap |\psi| \neq \emptyset$, contradicting the minimality of C_w. Therefore, w is a minimal world satisfying $\varphi \wedge \psi$ and since $w \leq_x y$, we get that $\mathfrak{M}, x \vDash [\varphi]\psi$. $\square$

15

We are now ready for our main theorem.

Theorem 4.6. *Broccoli logic = MCL.*

Proof. To show that *MCL* is *BL*, we need to show 1) that all axioms of section 4.1 are valid in *BL*, whose semantics was given in section 1 and 2) that if a principle is not derivable in *MCL*, then there is a broccoli countermodel.

Showing that the *MCL* axioms are valid in the *BL*-models of section 1 is straightforward. We show that axiom (5) is valid and leave the others to the reader. Let $\mathfrak{M}$ be an arbitrary broccoli model and let $u \in U$ be arbitrary. If $\neg\langle\varphi]\top \notin u$, i.e., if there is no revision by φ, then the thesis is vacuously true. Hence, assume that there is a revision by φ. Assume furthermore that $\mathfrak{M}, u \vDash [\varphi]\psi \wedge [\varphi]\theta$. Since $\mathfrak{M}, u \vDash [\varphi]\psi, |\varphi| \cap |\psi| \neq \emptyset$. Let $Z\mu(\mathcal{B} \bullet |\varphi \wedge \psi|)$ be a minimal set of $\mathcal{B}$ intersecting $|\varphi \wedge \psi|$. Then for every $z \in Z, x \in |\varphi| \cap |\psi|$ implies that $z \in |\varphi| \subseteq |\theta|$. Hence, $\mathfrak{M}, u \vDash [\varphi \wedge \psi]\theta$.

To show that if a principle is not provable in *MCL*, then there is a broccoli countermodel to φ, we use the completeness result of Burgess. If $MCL \nvdash \varphi$ for some φ, then there is a finite model $\mathfrak{M} = (U, R, V)$ and a world $u \in U$ such that $\mathfrak{M}, u \nvDash \varphi$. [5] By lemma 4.5, $BROC(\mathfrak{M}), u \nvDash \varphi$. Therefore, $BROC(\mathfrak{M})$ is a broccoli countermodel to φ. This completes the proof of theorem 4.6. □

Corollary 4.7. *BL is decidable.*

5   Conclusion

Our goal was to generalize onion semantics to capture relational belief revision; the result was *BL*. It turns out that *BL* is equivalent to a well-known conditional logic, the Burgess-Veltman minimal conditional logic. This is a fortunate outcome, as it saves a lot of work in coming up with a completeness result expanding on the minimal revisional logic of section 2. The major difficulty along the latter line was to devise an appropriate generalized arrow condition yielding generalized selection functions, and this is still an open question. Another open question is the rôle played by the $[\varphi\rangle$ modality in *BL*: what is the complete minimal logic of $[\varphi]\psi$ and $[\varphi\rangle\psi$ over the Burgess-Veltman models? An advantage of *MCL* over *BL* is that it avoids the problem of choosing an appropriate generalized limit condition by dropping the sphere representation altogether. A lesson should be drawn here, namely that, as so often over the past years, we see that logics of belief revision are largely conditional logics.

---

[5] Burgess proves that *MCL* has the finite model property.

References

[1] C. Alchourrón, P. Gärdenfors, and D. Makinson, *On the logic of theory change: partial meet contraction and revision fucntions*, The journal of symbolic logic **50** (1985), 510–530.

[2] P. Blackburn, M. de Rijke, and Y. Venema, *Modal logic*, Cambridge tracts in theoretical computer science, no. 53, Cambridge University Press, 2001.

[3] J. P. Burgess, *Quick completeness proofs for some logics of conditionals*, Notre Dame Journal of Fromal Logic **22** (1981), no. 1, 76–84.

[4] B. F. Chellas, *Modal logic: an introduction*, Cambridge University Press, 1980.

[5] A. Grove, *Two modellings for theory change*, Journal of philosophical logic **17** (1988), 157–170.

[6] D. Lewis, *Counterfactuals*, Harvard University Press, 1973.

[7] S. Lindström and W. Rabinowicz, *Epistemic entrenchment with incomparabilities and relational belief revision*, The Logic of Theory Change: Proc. of the Workshop (A. Fuhrmann and M. Morreau, eds.), Springer, Berlin, Heidelberg, 1991, pp. 93–126.

[8] D. Nute, *Conditional logic*, Handbook of Philosophical Logic (D. Gabbay and F. Guenthner, eds.), vol. II, D. Reidel Publishing Company, 1984, pp. 387–439.

[9] K. Segerberg, *Belief revision from the point of view of doxastic logic*, Bulleting of the IGPL **3** (1995), no. 4, 535–553.

[10] ———, *Proposal for a theory of belief revision along the lines of lindström and rabinowicz*, Fundamenta Informaticae **32** (1997), 183–191.

[11] ———, *Irrevocable belief revision in dynamic doxastic logic*, Notre Dame Journal of Formal Logic **39** (1998), no. 3, 287–306.

[12] ———, *The basic dynamic doxastic logic of AGM*, Frontiers in Belief Revision (M.-A. Williams and H. Rott, eds.), Kluwer Acadecmic Publishers, 2001, pp. 57–84.

[13] K. Segerberg and S. Lindström, *Modal logic and philosophy*, Handbook of modal logic, Elsevier.

[14] F. Veltman, *Logic for conditionals*, Ph.D. thesis, Department of Philosophy, University of Amsterdam, 1985.