Modal, Fuzzy, . . . , Vanilla Fixpoint Theories of Truth: A Uniform Approach

Melvin Fitting

Abstract Kripke's work on modal logic has been immensely influential. It hardly needs remarking that this is not his only work. Here we address his pioneering applications of fixpoint constructions to the theory of truth, and related work by others. In his fundamental paper on this he explicitly described a modal version, applying a fixpoint construction world by world within a modal frame. This can certainly be carried out, and doubtless has been somewhere. Others have suggested a variety of other extensions such as using the unit interval as the underlying space of truth values, or using a four valued logic instead of three, or various combinations of these. When many similar formal constructions have been proposed, one naturally asks what is the common core. Is there some setting in which things can be proved once and for all, with the various specific proposals seen as applications of this common core. In fact this is the case, with bilattices providing the desired structure. Much of such a development has already appeared in some form or other. It is the purpose of this article to bring everything together, and also add a few things. We present general results that, more or less, have everything currently in the fixpoint literature as special cases. Fortunately this does not make things more complicated, since the underlying proofs are essentially the same as they have always been. It is just that they are being carried out in generality rather than in specificity.

1 Introduction

In some areas of philosophy it is not uncommon to call on mathematical tools to help elucidate things. Saul Kripke is a master at this. Possible world semantics and fixed point theories of truth are two famous examples. They do not really solve philosophical problems themselves but, once having been introduced, they set the

---

Melvin Fitting
Graduate Center, City University of New York, Department of Philosophy (emeritus), 365 Fifth Avenue, New York City e-mail: melvin.fitting@gmail.com

language and the conceptual structures we all use when these problems are debated. When philosophical disagreements about modality or about a truth predicate arise these mathematical constructions generally serve the role of ground bass above which discussions play themselves out. Personally, I have always had an interest in tools as such, and that is my subject now. This is most decidedly a technical paper.

The famous paper on truth, Kripke (1975), introduced a kind of *truth revision operator* with the idea that fixed points of the operator were natural candidates for ways an “is true” predicate might behave. Roughly, fixed points are interpretations of a truth predicate that can not be bettered using the information they contain. A similar idea appeared, essentially simultaneously, in Martin and Woodruff (1975). This paper built on Kleene’s weak three valued logic, and used what are called *maximal* fixed points. Kripke took a broader view and examined three different logics: Kleene’s strong three valued logic, Kleene’s weak three valued logic, and supervaluations. In addition he investigated a range of different kinds of fixed points, with particular emphasis on the *least* one, but with the entire family playing a significant role. Both 1975 papers thought of classical logic as being at the heart of what they were doing, but Kripke noted that further extensions could be made.

> The present approach can be applied to languages containing modal operators. In this case we do not merely consider truth, but we are given, in the usual style of modal model theory, a system of possible worlds, and evaluate truth and $T(x)$ in each possible world. The inductive definition of the languages $\mathcal{L}_\alpha$ approximating to the minimal fixed point must be modified accordingly. We cannot give details here. Kripke (1975, pp. 712-713)

The choice of a three valued logic, while natural, has some technical drawbacks. The underlying truth value structure is not that of a complete lattice, but rather of a complete partial ordering. This makes the mathematics somewhat more complicated—detailed machinery needed for the Kripke version can be found in Fitting (1986). Visser (1984) built on the idea that the Belnap-Dunn four-valued semantics for first degree entailment was an alternate, and better, setting for a fixed point approach. Here the structure is a complete lattice, and this allows the use of the remarkably simple Knaster-Tarski fixed point theorem. In addition the four-valued semantics includes within its structure both Kleene’s strong three-valued logic and Graham Priest’s logic of paradox, LP, which is also suitable for a Kripke-style development. There have also been fixed point versions in which the underlying truth value space was fuzzy. And this does not exhaust the list of similar developments.

In every variation or extension of the 1975 papers, results *similar* to those have been shown, by *similar* methods. Whenever a mathematician sees formal structures with similarities, about which similar things are established by similar methods, that mathematician will ask, “what is the common ground of this similarity?” That common ground should be defined as an object of interest in its own right, and all these similar results should be proven just once, with all the rest as instances. For fixed point theories of truth this common ground is the *bilattice*. Why bilattices?

Bilattices are structures with two interconnected ordering relations. One of the relations can be used to give interpretations for logical connectives and quantifiers. These interpretations turn out to be monotonic with respect to the other order relation, and this fact can be used to show the existence fixed points for truth revision operators.

Kripke's work essentially took place in the simplest such example. He had only three truth values, roughly, true, false, and no-value-yet-assigned. This is actually not quite a full bilattice but it is the "consistent" part of the simplest one. In his setting one of the ordering relations was simply: false is less than true, with no-value-yet-assigned in between. Roughly speaking, this degree of truth ordering can be described by saying that an overall increase is either a decrease of falsehood or an increase of truth. The second ordering relation, the one for which monotonicity results obtain, is that no-value-yet-assigned is below both false and true, in effect a degree of information ordering.

Bilattices generalize the degree of truth and the degree of information orderings mentioned above. It will be seen that the bilattice family includes many natural examples that allow us to bring in modal logic, fuzzy logic, and other logics of interest as particular applications of general results. Some years ago I made use of bilattices in exactly this way in a series of papers: Fitting (1989, 1991, 1993, 1997, 2006), but things were rather spread out. The purpose of the present paper is to give a unified presentation, in full generality, with the items just mentioned as special cases, and with further applications provided as well.

The general structure of this paper is as follows.

In Sections 2 through 4 the essential facts about bilattices are presented, important examples are given, and methods of constructing them are discussed. It is shown how propositional logic, first order logic, and modal logics can be semantically interpreted using bilattices as many valued truth value spaces.

In Section 5 several fixed point theorems are proved, all following from the Knaster-Tarski theorem. This is not quite the fixed point theorem used by Kripke, since his underlying algebraic structures were not complete lattices while ours are. Moving to complete lattices has the advantage of simplifying the mathematics, while still allowing the construction to proceed more or less as Kripke did it.

Section 6 adds a truth predicate to the language of logic we are using. The fundamental monotonicity theorems are then proved.

Kripke worked with three valued logics. These can be thought of as based on the consistent values of the four-valued Belnap-Dunn structure. In Section 7 the notion of consistent value is generalized to bilattices, along with the notions of exact (a generalization of the classical truth values), and anticonsistent (dual to that of consistent).

Section 8 finally brings all the preceding together, extending Kripke's treatments of fixed point truth predicates using generalizations of Kleene's strong and Kleene's weak three valued logics, and generalizations of supervaluations, as well as a brief discussion involving asymmetric logics.

Section 9 examines some more specialized fixed point work, generalizing the notions of intrinsic and maximal. It also brings in so-called alternating fixed points, which actually first arose in computer science, but which have significance here as well.

2 Bilattices

There is an important feature common to fixpoint theories of truth generally: two partial orderings are involved. There is a space of truth values, with its ordering on *degree of truth* (or falsehood). There is also some truth revision operator involved, for which one wants fixed points, so some notion of approximation to a fixed point is desired. This always involves an ordering on what might be called the *degree of information*. The structure common to fixpoint truth theories, then, has two partial orderings, truth and information, somehow interconnected—and this is a *bilattice*.

Bilattices can have various conditions imposed. What we need here are bilattices with quite a bit of structure. They should be *complete, infinitarily interlaced bilattices, with a negation and a conflation that commute*. This is a lot of terminology, but the underlying ideas are all quite natural and are explained below. But as background we need lattices, bounded, and complete, so let us start there.

Definition 2.1 (Lattices) *A* partially ordered *set is a a set with a binary relation, $\le$, that is reflexive ($x \le x$), transitive ($x \le y$ and $y \le z$ imply $x \le z$), and anti-symmetric ($x \le y$ and $y \le x$ imply $x = y$).*

A lattice *is a partial ordering in which every pair of elements has a unique least upper bound (join) and a unique greatest lower bound (meet). (It follows that least upper and greatest lower bounds exist for every finite, non-empty set).*

A bounded *lattice is a lattice with a unique largest and a unique smallest element, usually called* top *and* bottom*.*

A lattice is complete *if every set (not just the finite ones) has a unique least upper bound and a unique greatest lower bound. (This implies boundedness, since the greatest lower bound of the empty set is top, and the least upper bound of the empty set is bottom.)*

Now we move to our real subject, bilattices. A pre-bilattice is simply a structure with two orderings. Once conditions connecting the two orderings are postulated, it is standard to drop the 'pre' qualification.

Definition 2.2 (Bilattices) *Terminology is as follows.*

Pre-Bilattice: *A pre-bilattice is a structure $\mathcal{B} = \langle \mathcal{B}, \le_t, \le_k \rangle$ with two lattice orderings; informally $\le_t$ is on degree of truth and $\le_k$ is on degree of information.*

Extreme Elements: *It is common to assume each lattice ordering is bounded, with a top and a bottom. We do so here. The smallest and largest members with respect to the truth ordering $\le_t$ will be denoted* **f** *and* **t***, and with respect to the information ordering $\le_k$, $\bot$ and $\top$.*

Bilattice Operations: *For $x, y \in \mathcal{B}$ their meet with respect to $\le_t$ is denoted $x \wedge y$ ("$\wedge$" is simply read "and"), the join is denoted $x \vee y$ ("$\vee$" is read "or"). With respect to $\le_k$ the meet is $x \otimes y$ ("$\otimes$" is read "consensus"), and join is $x \oplus y$ ("$\oplus$" is read "gullability" or "accept all").*

Interlaced: *It follows from the definitions of least upper and greatest lower bounds, that each of the four bilattice operations is monotone with respect to its ordering.*

For example, if $x \leq_k x'$ *and* $y \leq_k y'$ *then* $x \otimes y \leq_k x' \otimes y'$. *A bilattice is* interlaced *if each operation is monotone with respect to* both *orderings. As one example,* $x \leq_t x'$ *and* $y \leq_t y'$ *then* $x \otimes y \leq_t x' \otimes y'$.

Negation: *A bilattice has a* negation *if there is an involution* $\neg : \mathcal{B} \rightarrow \mathcal{B}$ *that reverses the truth ordering and leaves intact the information ordering. Thus* $\neg\neg x = x$ *and* $x \leq_t y$ *implies* $\neg y \leq_t \neg x$ *and* $x \leq_k y$ *implies* $\neg x \leq_k \neg y$.

Conflation: *A bilattice has a* conflation *if there is an involution* $- : \mathcal{B} \rightarrow \mathcal{B}$ *that reverses the information ordering and leaves intact the truth ordering. Thus* $--x = x$ *and* $x \leq_k y$ *implies* $-y \leq_k -x$ *and* $x \leq_t y$ *implies* $-x \leq_t -y$. *Negation and conflation commute if always* $\neg - x = -\neg x$.

Completeness: *A bilattice is* complete *if both* $\langle \mathcal{B}, \leq_t \rangle$ *and* $\langle \mathcal{B}, \leq_k \rangle$ *are complete lattices. We write* $\bigwedge S$ *and* $\bigvee S$ *for the meet and join of set* S *with respect to the* $\leq_t$ *ordering, and* $\prod S$ *and* $\sum S$ *for meet and join with respect to the* $\leq_k$ *ordering.*

Infinitarily Interlaced: *A bilattice is* infinitarily interlaced *if all meets and joins with respect to each ordering are monotone with respect to the other ordering, as well as to their own. As one example, suppose we have an indexing set* I, *both* $\{a_i \mid i \in I\}$ *and* $\{b_i \mid i \in I\}$ *are subsets of* $\mathcal{B}$, *and* $a_i \leq_t b_i$ *for each* $i \in I$. *Then* $\prod_i a_i \leq_t \prod_i b_i$.

The conditions above are not all independent. Completeness implies the existence of tops and bottoms for both orderings, and infinitarily interlaced implies interlaced. For applications in other areas not all of the conditions may be required, but we will want all of them except sometimes conflation. It is a lot to keep saying that we have a pre-bilattice that is interlaced, has a negation and a conflation that commute, is complete, and is infinitarily interlaced. Consequently we specialize terminology here. When reading other papers, please check usage there, to avoid confusion.

Terminology Convention 2.3 From now on a bilattice meeting all the conditions above, omitting conflation, will simply be referred to here as a *bilattice*. This is an abuse of terminology, to be sure, but it simplifies verbiage and should cause no general problems. By a *bilattice with conflation* we mean a bilattice in which there is a conflation operation that commutes with negation.

Thinking of a bilattice as a space of truth values, logical connectives for conjunction, disjunction, and negation can be interpreted using the truth operations of the bilattice. It should be no problem if we use "$\wedge$" for both the logical connective and for the bilattice operation, since context will sort things out. Similarly for "$\vee$" and "$\neg$". Quantifiers will be interpreted using the infinitary meet, "$\bigwedge$", and join, "$\bigvee$", operations of the truth ordering. The information operations will be used too, but for purposes other than directly interpreting logical operations.

Example 2.4 Figure 1 shows an example of a bilattice with conflation, the simplest one since it has only the four extreme values which every bilattice must have. It is certainly the most important example, and we refer to Arieli and Avron (1998) for some of the reasons why. It is standard to display bilattices using a double Hasse diagram. The truth ordering is from left to right, and the information ordering is from bottom upwards. The bilattice $\mathcal{FOUR}$ derives from Belnap and Dunn, see Belnap (1977), and its truth ordering provides the standard semantics for first degree entailment (which will not play a role here). It is actually an example of a *distributive* bilattice, that is, each operation distributes over all the others. For instance, $x \wedge (y \oplus z) = (x \wedge y) \oplus (x \wedge z)$. It is not hard to show that distributivity implies interlacing. The converse is not true. We will not mention distributivity further since interlacing has turned out to be the more important concept for bilattices. Since $\mathcal{FOUR}$ is interlaced it is trivially infinitarily interlaced, being finite. Negation is left-right symmetry. Conflation is vertical symmetry. We will revisit this example from time to time.

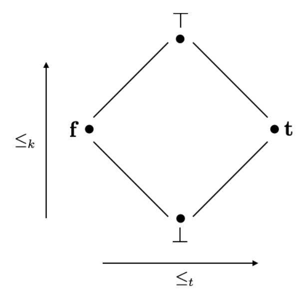

Fig. 1: The Bilattice $\mathcal{FOUR}$

Proposition 2.5 (De Morgan Laws) *In the following the left entries for are bilattices generally, and the right are for bilattices with conflation. For a set S of bilattice members we write $\neg S$ for $\{\neg x \mid x \in S\}$, and similarly for $-S$.*

$$\begin{array}{ll}
\neg(x \wedge y) = \neg x \vee \neg y & -(x \wedge y) = -x \wedge -y \\
\neg(x \vee y) = \neg x \wedge \neg y & -(x \vee y) = -x \vee -y \\
\neg(x \otimes y) = \neg x \otimes \neg y & -(x \otimes y) = -x \oplus -y \\
\neg(x \oplus y) = \neg x \oplus \neg y & -(x \oplus y) = -x \otimes -y \\
\neg \bigwedge S = \bigvee \neg S & - \bigwedge S = \bigwedge -S \\
\neg \bigvee S = \bigwedge \neg S & - \bigvee S = \bigvee -S \\
\neg \prod S = \prod \neg S & - \prod S = \sum -S \\
\neg \sum S = \sum \neg S & - \sum S = \prod -S
\end{array}$$

The following says the four extreme members of a bilattice are always tightly connected. It can be proved directly but here it is a special case of Proposition 5.8.

Proposition 2.6 *In any bilattice:* $\mathbf{t} \otimes \mathbf{f} = \bot$, $\mathbf{t} \oplus \mathbf{f} = \top$, $\top \wedge \bot = \mathbf{f}$, *and* $\top \vee \bot = \mathbf{t}$.

3 Constructing Bilattices

There is a universal way of constructing bilattices, and also a function space construction that is not universal but is essential for what we will do. This does not exhaust the subject of bilattice constructions, but is sufficient for our needs here.

3.1 Bilattice Products

We are about to present an intuitive and universal construction for bilattices. It entered the bilattice community, with various names, through Ginsberg (1988), Fitting (1990, 1991), Avron (1996), Pynko (2000). It turns out that, in fact, it was previously known using different terminology and in a different context, see Gargov (1999), Davey (2013). We are primarily interested in a special case here, because of our Terminology Convention 2.3, but we do mention more general versions. In the following, note the reversal involving the second component, for the $\leq_t$ ordering.

Definition 3.1 (Bilattice Product) *Let* $L_1 = \langle L_1, \leq_1 \rangle$ *and* $L_2 = \langle L_2\ \leq_2 \rangle$ *be two lattices. The* bilattice product*,* $L_1 \odot L_2$*, is the structure* $\langle L_1 \times L_2, \leq_t, \leq_k \rangle$ *where*

1. $\langle x_1, y_1 \rangle \leq_t \langle x_2, y_2 \rangle$ *if* $x_1 \leq_1 x_2$ *and* $y_2 \leq_2 y_1$*,*
2. $\langle x_1, y_1 \rangle \leq_k \langle x_2, y_2 \rangle$ *if* $x_1 \leq_1 x_2$ *and* $y_1 \leq_2 y_2$*.*

There is a natural intuition behind this. Think of L_1 as representing degree of evidence *for* some proposition, and L_2 as representing degree of evidence *against*. These need not be measured the same way. A value in a bilattice product separately records evidence for and evidence against. Degree of information (the $\leq_k$ ordering) goes up if we have the same or more evidence both for and against. Degree of truth (the $\leq_t$ ordering) goes up if evidence for goes up while evidence against goes down.

The bilattice product construction is quite versatile. In the following discussion we use $\sqcap_1$ and $\sqcup_1$ for meet and join of lattice L_1 and $\sqcap_2$ and $\sqcup_2$ for meet and join of lattice L_2. Likewise if the lattices are bounded, we use 0_1 and 1_1 for bottom and top of L_1, and similarly for L_2. If the lattices are complete we use $\bigsqcap_1$ and $\bigsqcup_1$ for the meet and join of arbitrary subsets of L_1 and similarly for L_2.

1. $L_1 \odot L_2$ is a lattice with respect to both orderings, with the following operations.

$$\begin{aligned}
\langle x_1, y_1 \rangle \oplus \langle x_2, y_2 \rangle &= \langle x_1 \sqcup_1 x_2, y_1 \sqcup_2 y_2 \rangle \\
\langle x_1, y_1 \rangle \otimes \langle x_2, y_2 \rangle &= \langle x_1 \sqcap_1 x_2, y_1 \sqcap_2 y_2 \rangle \\
\langle x_1, y_1 \rangle \vee \langle x_2, y_2 \rangle &= \langle x_1 \sqcup_1 x_2, y_1 \sqcap_2 y_2 \rangle \\
\langle x_1, y_1 \rangle \wedge \langle x_2, y_2 \rangle &= \langle x_1 \sqcap_1 x_2, y_1 \sqcup_2 y_2 \rangle
\end{aligned}$$

2. $L_1 \odot L_2$ always meets the interlacing conditions.
3. If L_1 and L_2 are bounded lattices $L_1 \odot L_2$ has extreme elements: $\langle 0_1, 0_2 \rangle$ and $\langle 1_1, 1_2 \rangle$ for the information ordering; $\langle 0_1, 1_2 \rangle$ and $\langle 1_1, 0_2 \rangle$ for the truth ordering.

4. If L_1 and L_2 are complete lattices then $L_1 \odot L_2$ is a complete bilattice, with the following infinitary operations.

$$\sum S = \langle \bigsqcup_1 \{x \mid \langle x, y\rangle \in S \text{ where } y \in L_2\}, \bigsqcup_2 \{y \mid \langle x, y\rangle \in S \text{ where } x \in L_1\}\rangle$$
$$\prod S = \langle \bigsqcap_1 \{x \mid \langle x, y\rangle \in S \text{ where } y \in L_2\}, \bigsqcap_2 \{y \mid \langle x, y\rangle \in S \text{ where } x \in L_1\}\rangle$$
$$\bigvee S = \langle \bigsqcup_1 \{x \mid \langle x, y\rangle \in S \text{ where } y \in L_2\}, \bigsqcap_2 \{y \mid \langle x, y\rangle \in S \text{ where } x \in L_1\}\rangle$$
$$\bigwedge S = \langle \bigsqcap_1 \{x \mid \langle x, y\rangle \in S \text{ where } y \in L_2\}, \bigsqcup_2 \{y \mid \langle x, y\rangle \in S \text{ where } x \in L_1\}\rangle$$

5. If L_1 and L_2 are complete lattices then $L_1 \odot L_2$ is infinitarily interlaced.
6. If $L_1 = L_2 = L$, that is, if the two lattices are the same, then $L \odot L$ has a negation, $\neg\langle x, y\rangle = \langle y, x\rangle$.

Conflation needs some additional machinery.

Definition 3.2 (De Morgan Lattice) *A De Morgan lattice is a bounded distributive lattice (Definition 2.1) with a De Morgan involution operation $x \mapsto \overline{x}$ that is order reversing, $x \leq y$ implies $\overline{y} \leq \overline{x}$, and is an involution, $\overline{\overline{x}} = x$. Since we will never need distributivity, we call a structure a* non-distributive *De Morgan lattice if it meets all the conditions except possibly for the distributive laws.*

It is easy to show De Morgan laws hold in any non-distributive De Morgan lattice. One more item is now added to the list above.

7. If $L_1 = L_2 = L$ is a non-distributive De Morgan lattice, then $L \odot L$ has a conflation, $-\langle x, y\rangle = \langle \overline{y}, \overline{x}\rangle$, and it commutes with negation, item 6.

We now have a simple way of constructing bilattices. It is completely general. For instance, consider items 2 and 3. In the converse direction it can be shown that every bilattice with extreme elements that meets the interlacing conditions is isomorphic to $L_1 \odot L_2$ where both L_1 and L_2 are bounded lattices. And so on. Thus we actually have full representation theorems.

3.2 Function Space Bilattices

This is another way of constructing bilattices that will be of much use to us. It is not universal, unlike bilattice product, but it does help to unify several constructions that have appeared in the literature on fixed point theories of truth.

Definition 3.3 (Function Space Bilattice) *Let $\mathcal{B} = \langle \mathcal{B}, \leq_t, \leq_k\rangle$ be a bilattice, and $\mathcal{A}$ be a non-empty set. The function space $\mathcal{B}^{\mathcal{A}} = \langle \mathcal{B}^{\mathcal{A}}, \leq_t, \leq_k\rangle$ is a bilattice in which we take as domain the set consisting of all functions from $\mathcal{A}$ to $\mathcal{B}$, and with orderings that are pointwise, which means that for $f, g \in \mathcal{B}^{\mathcal{A}}$:*

1. $f \leq_t g$ iff $f(x) \leq_t g(x)$ for all $x \in \mathcal{A}$, *2.* $f \le_k g$ *iff* $f(x) \le_k g(x)$ *for all* $x \in \mathcal{A}$.

Notice that we have used $\le_t$ and $\le_k$ to denote partial orderings in both $\mathcal{B}$ and $\mathcal{B}^{\mathcal{A}}$. This keeps notation relatively simple, and context should always clarify which bilattice is meant. A function space bilattice inherits properties from the underlying bilattice. We leave verifications of the following to you.

1. $\mathcal{B}^{\mathcal{A}}$ is a lattice with respect to both orderings. For instance, the meet with respect to $\le_k$ is easily verified to be the function given by $(f \oplus g)(x) = f(x) \oplus g(x)$, and similarly with the other meet and both joins.
2. $\mathcal{B}^{\mathcal{A}}$ is interlaced (since $\mathcal{B}$ is).
3. $\mathcal{B}^{\mathcal{A}}$ is bounded with respect to both orderings. For instance, the least member of $\mathcal{B}^{\mathcal{A}}$ in the pointwise $\le_k$ ordering is the function that is identically $\bot$ on $\mathcal{A}$, where $\bot$ is the least member of $\mathcal{B}$ in the $\le_k$ ordering on $\mathcal{B}$. Similarly for the other three extreme elements.
4. $\mathcal{B}^{\mathcal{A}}$ is a complete bilattice (since $\mathcal{B}$ is). If $S \subseteq \mathcal{B}^{\mathcal{A}}$, $\sum S$ is the function such that $(\sum S)(x) = \sum\{f(x) \mid f \in S\}$, and similarly for the other cases.
5. $\mathcal{B}^{\mathcal{A}}$ is infinitarily interlaced.
6. $\mathcal{B}^{\mathcal{A}}$ has a negation,and $(\neg f)(x) = \neg(f(x))$.
7. If $\mathcal{B}$ has a conflation that commutes with negation, the same is true of $\mathcal{B}^{\mathcal{A}}$,and $(-f)(x) = -(f(x))$.

If we have a space of truth values that constitutes a bilattice, such as $\mathcal{FOUR}$, the space of valuations in this bilattice, which is a function space, is actually a bilattice. It is here that most of the fixpoint machinery is actually used.

4 Bilattice Examples

We have already seen the most basic bilattice example, $\mathcal{FOUR}$, in Figure 1, which is an easy bilattice product construction. Start with the simplest Boolean algebra with underlying space $\{0, 1\}$, and with $0 \le 1$. Identifying 0 with falsehood and 1 with truth, the meet and join are those of classical propositional logic. There is a natural De Morgan involution, switch 0 and 1—this corresponds to classical negation. The bilattice product of this structure with itself is, isomorphically, $\mathcal{FOUR}$, with $\bot$ being $\langle 0, 0\rangle$, $\top$ being $\langle 1, 1\rangle$, **f** as $\langle 0, 1\rangle$ and **t** as $\langle 1, 0\rangle$. We now move on to several less familiar examples.

Example 4.1 (Bilattices of Sets of Agents) Suppose we have two sets of agents, L_1 and L_2. Members of L_1 get to vote that a particular informal proof of some proposition is correct, or else offer no opinion. Members of L_2 play a similar role with respect to the correctness of some informally presented countermodel to the proposition. These are our 'for' and 'against' agents. The two kinds of agents could even be drawn from different populations. A natural lattice ordering relation is simply subset, so degree goes up if more agents join a group. We have no requirement that everybody speaks up on every proposition; indeed there might be no agents either for or against some proposition. We have no requirement that agents be consistent, some might belong to both groups and decide to accept both a proof and a countermodel. This can easily happen since both proofs and countermodels might be allowed to be informal, so there is the possibility of error or misunderstanding. At any rate, a member of the bilattice product $L_1 \odot L_2$ encodes what set of agents is for, and what set is against.

Degree of information increases if more agents join one or another (or both) of the sets. Degree of truth goes up if more join the set for or leave the set against. The bottom in the truth ordering is $\langle \emptyset, L_2 \rangle$, nobody accepts an informal proof, everybody accepts a countermodel. Similarly for the other extreme bilattice values.

If the sets of agents, for and against, are drawn from the same population, that is if $L_1 = L_2$, then negation simply reverses the positions of those who are for and those who are against. Further, if $L_1 = L_2$ the appropriate De Morgan involution is relative compliment with respect to the entire set of agents. Then conflation in the bilattice product of agents replaces the set of agents for something with the set of those who were not against it, and the set against with the set of those who were not for.

$\mathcal{FOUR}$ can be thought of as an agent bilattice with a single agent. There are two agent subsets, none and all, and these can be identified with the members of $\{0, 1\}$.

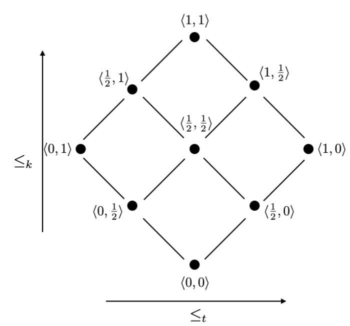

Fig. 2: The Bilattice $\mathcal{NINE}$

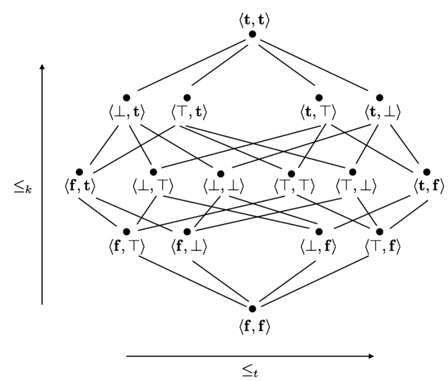

Fig. 3: The Bilattice $\mathcal{SIXTEEN}$

Example 4.2 (Bilattice $\mathcal{NINE}$) This example is similar to $\mathcal{FOUR}$, but now start with the space $\{0, \frac{1}{2}, 1\}$ with the standard numerical ordering. The bilattice product of this with itself is in Figure 2. Properties are similar to those of $\mathcal{FOUR}$. It can be thought of as a biltice built on Kleene's strong three valued logic, or Priest's logic of paradox.

Example 4.3 (Bilattice $\mathcal{SIXTEEN}$) Start with the lattice from $\mathcal{FOUR}$ just using the $\leq_t$ ordering. The bilattice product of this with itself is shown in Figure 3. It has features of interest that will be discussed from time to time.

Example 4.4 (A Fuzzy Bilattice) Examples so far have been finite. All have been complete bilattices, simply because they are finite. Now we look at a 'fuzzy' example. Consider the closed unit real interval, $[0, 1]$, with the usual numerical ordering. Form the bilattice product of this with itself, so a bilattice truth value is a pair $\langle x, y\rangle$ where x can be thought of as a degree of belief in some proposition, and y as a corresponding degree of doubt. (Think of belief and doubt as being independently arrived at.) Loosely, one can think of it as $\mathcal{FOUR}$, but 'filled in'. This is a bilattice that is complete, not because it is finite, but because the unit interval ordering is a complete one. It is also a bilattice with conflation, where $-\langle x, y\rangle = \langle 1 - y, 1 - x\rangle$.

Example 4.5 (Propositional Valuations) The previous examples made use of bilattice products. Function space bilattices also arise quite naturally. Let $\mathcal{A}$ be a set of propositional atoms, from which formulas of propositional language $\mathcal{L}$ are built using connectives $\wedge$, $\vee$, and $\neg$. Let our space of truth values be a bilattice $\mathcal{B}$, for instance any of the examples discussed above. A *propositional valuation* is any mapping from the set of atoms $\mathcal{A}$, to the bilattice $\mathcal{B}$, an arbitrary member of $\mathcal{B}^{\mathcal{A}}$. This gives us a bilattice, $\langle \mathcal{B}^{\mathcal{A}}, \leq_t, \leq_k\rangle$, using pointwise orderings as described in Section 3.2. For propositional valuations v and w we have $v \leq_k w$ if $v(P) \leq_k w(P)$ for every propositional letter P, and similarly for $\leq_t$.

One could also consider a function space bilattice $\mathcal{B}^{\mathcal{L}}$, where $\mathcal{L}$ is the entire set of propositional formulas, and not just the atomic ones. This is not of much interest as a whole, but it does have an important subspace. Each propositional valuation extends to all formulas in a unique way so that the following conditions are met, using the truth operations from the bilattice $\mathcal{B}$ in place of conventional truth tables.

$$
\begin{aligned}
v(X \wedge Y) &= v(X) \wedge v(Y)\\
v(X \vee Y) &= v(X) \vee v(Y)\\
v(\neg X) &= \neg v(X).
\end{aligned}
$$

On the left we have operation symbols of the language $\mathcal{L}$, and on the right operations from the truth ordering in the bilattice $\mathcal{B}$. Note that we have used the same notation for valuations in $\mathcal{B}^{\mathcal{A}}$ and for their extensions in $\mathcal{B}^{\mathcal{L}}$, relying on context to clarify which is meant. Valuations, extended as above to mappings on all formulas, constitute an obvious natural subspace of the bilattice $\mathcal{B}^{\mathcal{L}}$. Of the two bilattice orderings restricted to this subspace, $\leq_t$ is not interesting (because of the behavior of negation), but $\leq_k$ most decidedly is of interest.

Proposition 4.6 (Propositional Monotonicity) *For propositional valuations $v, w \in \mathcal{B}^{\mathcal{A}}$,*

$$v \leq_k w \text{ in } \mathcal{B}^{\mathcal{A}} \text{ if and only if } v \leq_k w \text{ in } \mathcal{B}^{\mathcal{L}}.$$

Stated more explicitly, for valuations v and w, extended to all formulas as above, $v(P) \leq_k w(P)$ for all propositional atoms P if and only if $v(X) \leq_k w(X)$ for all formulas X. The implication from right to left is trivial. The implication from left to right is shown by a simple induction on formula complexity, and we leave it to you.

Example 4.7 (Quantificational Valuations) Let formal language $\mathcal{L}$ now be first-order, with relation symbols but no constant or function symbols (for convenience only). Assume language $\mathcal{L}$ has both universal and existential quantifiers. Let $\mathcal{D}$ be some non-empty domain, which we think of quantifiers as ranging over. Doing things a little differently from the standard (but following Smullyan (1968)) rather than allowing free variables to occur in formulas and working with a function assigning members of the domain to those free variables, we simply allow members of the domain to occur directly in formulas. By a $\mathcal{D}$ sentence we mean a formula of $\mathcal{L}$ in which all free variable occurrences have been replaced with members of $\mathcal{D}$—we write $\mathcal{L}(\mathcal{D})$ for the set of all $\mathcal{D}$ sentences. Atomic $\mathcal{D}$ sentences are of the form $P(a_1, \ldots, a_n)$ where P is a relation symbol and $a_1, \ldots, a_n$ are members of $\mathcal{D}$. We call the set of atomic $\mathcal{D}$ sentences $\mathcal{A}(\mathcal{D})$.

Let $\mathcal{B}$ be a bilattice of intended truth values. A *quantificational valuation* in $\mathcal{B}$ is a mapping from atomic $\mathcal{D}$ sentences to members of $\mathcal{B}$, a member of function space bilattice $\mathcal{B}^{\mathcal{A}(\mathcal{D})}$. Quantificational valuation v can be extended to all $\mathcal{D}$ sentences, that is, to a mapping in $\mathcal{B}^{\mathcal{L}(\mathcal{D})}$. Propositional connectives are handled exactly as in Example 4.5. The quantificational cases are as follows.

$$v(\forall x F(x)) = \bigwedge \{v(F(d)) \mid d \in \mathcal{D}\}$$
$$v(\exists x F(x)) = \bigvee \{v(F(d)) \mid d \in \mathcal{D}\}$$

Proposition 4.6 extends to the following, whose proof we again leave to you.

Proposition 4.8 (Quantificational Monotonicity) *For quantificational valuations* $v, w \in \mathcal{B}^{\mathcal{A}(\mathcal{D})}$,

$$v \leq_k w \textit{ in } \mathcal{B}^{\mathcal{A}(\mathcal{D})} \textit{ if and only if } v \leq_k w \textit{ in } \mathcal{B}^{\mathcal{L}(\mathcal{D})}.$$

Example 4.9 (Modal Valuations) Both propositional and quantified languages extend to include modal operators. We only discuss the quantified version, and we assume we have constant domain models. Modifying the work to allow varying domain is not hard, but the essential ideas are more easily discerned in the present, simpler, setting. We start with the first-order machinery of Example 4.7, a bilattice $\mathcal{B}$, a domain $\mathcal{D}$ of quantification, and a set $\mathcal{A}(\mathcal{D})$ of atomic $\mathcal{D}$ sentences incorporating members of $\mathcal{D}$ directly. The first-order language is extended with modal operators $\Box$ and $\Diamond$ in the usual way. We continue to use $\mathcal{L}(\mathcal{D})$ for the set of $\mathcal{D}$ sentences, but now allowing modal operators. As new semantic machinery suppose we have a modal *frame*, $\mathcal{F} = \langle \mathcal{G}, \mathcal{R} \rangle$, where $\mathcal{G}$ is a non-empty set of possible worlds and $\mathcal{R}$ is a binary accessibility relation. We assume the basics of possible world semantics for modal logics are understood.

Form the function space bilattice $\mathcal{B}^{\mathcal{G}}$, with pointwise orderings. Since $\mathcal{B}^{\mathcal{G}}$ is a bilattice, $(\mathcal{B}^{\mathcal{G}})^{\mathcal{A}(\mathcal{D})}$ is also a bilattice, of the kind discussed in Example 4.7 except that the underlying bilattice is now itself a function space bilattice. A member of this bilattice is a valuation function $v : \mathcal{A}(\mathcal{D}) \to (\mathcal{G} \to \mathcal{B})$ that assigns to each atomic $\mathcal{D}$ sentence, and each possible world, a truth value in the bilattice $\mathcal{B}$. We call members of this *quantificational modal valuations*.

Suppose $v \in (\mathcal{B}^{\mathcal{G}})^{\mathcal{A}(\mathcal{D})}$; in Example 4.5 we said how to extend v at the atomic level, adding cases involving propositional connectives, and this carries over to here. We sketch how it gives us the appropriate behavior for propositional connectives in the modal setting as well. Suppose that $v \in (\mathcal{B}^{\mathcal{G}})^{\mathcal{L}(\mathcal{D})}$ meets the conditions for $\wedge$, $\vee$, and $\neg$ from Example 4.5. Then $v(X \wedge Y) = v(X) \wedge v(Y)$. But v maps to $\mathcal{B}^{\mathcal{G}}$, so $v(X \wedge Y)$ is a function from $\mathcal{G}$ to $\mathcal{B}$, and similarly for $v(X)$ and $v(Y)$. Then for a possible world $\Gamma \in \mathcal{G}$, $v(X \wedge Y)(\Gamma) = (v(X) \wedge v(Y))(\Gamma)$. The $\wedge$ on the right is an operator of $\mathcal{B}^{\mathcal{G}}$, and behaves pointwise, so $(v(X) \wedge v(Y))(\Gamma) = v(X)(\Gamma) \wedge v(Y)(\Gamma)$. Thus we have $v(X \wedge Y)(\Gamma) = v(X)(\Gamma) \wedge v(Y)(\Gamma)$. This is the usual way conjunction is understood in a possible world model. Disjunction and negation are similar.

The quantificational cases are as in Example 4.7, and here too we inherit appropriate world by world behavior automatically. For instance, we have the following: $v(\forall x\varphi(x))(\Gamma) = \bigwedge\{v(\varphi(d))(\Gamma) \mid d \in \mathcal{D}\}$, and similarly for $\exists$.

What must be added is extension conditions to cover the modal operators. Recall, we are working with respect to a frame $\mathcal{F} = \langle \mathcal{G}, \mathcal{R} \rangle$. We now add the following, for each $\Gamma \in \mathcal{G}$, each $v \in (\mathcal{B}^{\mathcal{G}})^{\mathcal{A}(\mathcal{D})}$ and each $\mathcal{D}$ sentence X.

$$v(\Box X)(\Gamma) = \bigwedge\{v(X)(\Delta) \mid \Delta \in \mathcal{G} \text{ and } \Gamma\mathcal{R}\Delta\}$$
$$v(\Diamond X)(\Gamma) = \bigvee\{v(X)(\Delta) \mid \Delta \in \mathcal{G} \text{ and } \Gamma\mathcal{R}\Delta\}$$

We have covered all cases for extending a quantificational modal valuation $v \in (\mathcal{B}^{\mathcal{G}})^{\mathcal{A}(\mathcal{D})}$ to $v \in (\mathcal{B}^{\mathcal{G}})^{\mathcal{L}(\mathcal{D})}$, allowing modal operators. It would be a good exercise to show that $v(\neg\Box\neg X)(\Gamma) = v(\Diamond X)(\Gamma)$ for every $\mathcal{D}$ sentence X and every $\Gamma \in \mathcal{G}$.

Propositions 4.6 and 4.8 now easily extend to the following.

Proposition 4.10 (Modal Monotonicity) *Using a modal first-order language $\mathcal{L}$ and a frame $\mathcal{F} = \langle \mathcal{G}, \mathcal{R} \rangle$, for modal valuations $v, w \in (\mathcal{B}^{\mathcal{G}})^{\mathcal{A}(\mathcal{D})}$,*

$$v \leq_k w \text{ in } (\mathcal{B}^{\mathcal{G}})^{\mathcal{A}(\mathcal{D})} \text{ if and only if } v \leq_k w \text{ in } (\mathcal{B}^{\mathcal{G}})^{\mathcal{L}(\mathcal{D})}.$$

It is clear that a rich variety of bilattices is available. $\mathcal{FOUR}$ and its various subspaces is often the setting for fixed point theories of truth but as we will see, essentially all the work applies quite generally. We suggest that Example 4.1 involving sets of agents is worth considering as a possible setting for thinking about the mutual understanding of a truth predicate in a community. This might be particularly interesting when the language involves modalities, perhaps even multimodalities, with a knowledge modal operator for each agent. Combining modal models with fuzzy bilattices, Example 4.4, is worth investigating. We leave such things to others.

5 Some Fixed Point Theorems

The application of fixed point theorems to theories of truth originated in Kripke (1975) and Martin and Woodruff (1975), but details are quite different between the two papers. Kripke used a *least* fixed point construction while Martin and Woodruff relied on *maximal* fixed points. Neither worked in a full bilattice setting, but rather in a substructure of $\mathcal{FOUR}$ that we will be generalizing here. In a bilattice framework the simple and elegant fixed point result of Knaster and Tarski is the one that is most pertinent, and is what we will use here. This has, for its setting, a complete lattice, see Definition 2.1. In this section we begin with lattices, then move to bilattices. On the one hand we do not need the full Knaster-Tarski theorem, which has to do with the structure of the entire set of fixed points of a monotone mapping. All we need are greatest and least fixed points, and we skip the rest. On the other hand we will need both the original version and some simple generalizations.

Definition 5.1 *Let $\langle L, \leq\rangle$ be a lattice. A mapping $f : L \to L$ is* monotone *if it is order preserving, that is, if $x \leq y$ then $f(x) \leq f(y)$. A member $x \in L$ is: a* fixed point *of f if $f(x) = x$, a* pre-fixed point *if $f(x) \leq x$, and a* post-fixed point *if $x \leq f(x)$.*

Proposition 5.2 (Knaster-Tarski) *Let f be a monotone mapping on a complete lattice. Then:*

1. *f has a smallest and a greatest fixed point.*
2. *If a is a post-fixed point of f then there is a fixed point of f, least above a.*
3. *If b is a pre-fixed point of f then there is a fixed point of f, greatest below b.*

The Knaster-Tarski theorem has two quite different proofs. One is short, self-contained, and somewhat magical; the other is substantially longer but provides more intuition. It is not self-contained, making use of facts from set theory concerning ordinal numbers. We sketch the longer proof, summarizing things properly belonging to set theory. We begin with ordinal numbers.

Proof Sketch The sequence of ordinals starts with the natural numbers, $0, 1, 2, \ldots$, followed by the smallest infinite ordinal, ω, which is followed by $\omega + 1$, $\omega + 2$ and so on. After these come the second infinite ordinal, $\omega \cdot 2$, then $\omega \cdot 2 + 1$, $\omega \cdot 2 + 2$, and so on (a phrase that covers a lot). The ordinal ω is the first *limit* ordinal, $\omega \cdot 2$ is the second, with infinitely many limit ordinals beyond. Ordinals divide into three groups: successor ordinals, which can be written as $\alpha + 1$ where α is the immediate predecessor; limit ordinals, with predecessors but no immediate one; and 0, the only ordinal that is neither a limit ordinal nor a successor ordinal. The sequence of ordinals is not itself a set and so cannot be placed in a 1-1 correspondence with any set.

We begin with part 2. Let $\langle L, \leq\rangle$ be a complete lattice (assumed to be a set), with $f : L \to L$ a monotone function, and a a post-fixed point of f, that is $a \leq f(a)$. We sketch why there is a fixed point for f above a, least among fixed points above a.

Use transfinite recursion to define a mapping φ from the class of ordinals to L, call it an *ordinal sequence*.

$$\begin{aligned}
\varphi(0) &= a \text{ a post-fixed point of } f\\
\varphi(\alpha + 1) &= f(\varphi(\alpha)) \text{ for a successor ordinal } \alpha + 1\\
\varphi(\lambda) &= \bigvee\{\varphi(\alpha) \mid \alpha < \lambda\} \text{ for a limit ordinal } \lambda
\end{aligned}$$

The key fact about φ is that it is monotonic, in the sense that if α and β are ordinals and $\alpha \leq \beta$, then in lattice L we have $\varphi(\alpha) \leq \varphi(\beta)$. The proof of this is set theoretic and a bit technical so we omit it. It follows that the sequence $\varphi(\alpha)$ increases in L as ordinal α increases. That increase cannot be strict because otherwise we would be pairing the ordinals up in a 1-1 fashion with a subset of L, but the collection of ordinals cannot be paired up with the members of a set. Then at some point we must have $\varphi(\alpha) = \varphi(\alpha + 1)$. Let α_0 be the least ordinal for which this happens. Then $\varphi(\alpha_0)$ is a fixed point of f because $f(\varphi(\alpha_0)) = \varphi(\alpha_0 + 1) = \varphi(\alpha_0)$. Further, $\varphi(\alpha_0)$ is the *least* fixed point for f above a, because if F is any fixed point of f above a then $\varphi(\alpha) \leq F$ for all ordinals, so in particular, $\varphi(\alpha_0) \leq F$. This is shown by transfinite induction on α, and again we omit the set theoretic proof.

Part 3 has a similar proof, but with everything dualized.

Part 1 now follows since $\bot$ is a post-fixed point of f with every member of the lattice above it, and dually $\top$ is a pre-fixed point with every member below it. □

We not only have the existence of least and greatest fixed points in a complete lattice, but a way of proving that they have certain properties. Standard terminology: in a partial ordering two members x and y are called *comparable* if $x \leq y$ or $y \leq x$, and a subset is a *chain* if any two members are comparable.

Definition 5.3 *Let $\langle L, \leq \rangle$ be a complete lattice, f be a monotone function on L, a be a post-fixed point of f, and P be a subset of L. We say P is f*-inductive starting at *a if:*

1. *$a \in P$;*
2. *If $x \in P$ then $f(x) \in P$;*
3. *P is* closed under chain sups*, meaning that if $S \subseteq P$ and S is a chain then $\bigvee S \in P$.*

Corollary 5.4 *If f is a monotone function on complete lattice L, a is a post-fixed point of f, and $P \subseteq L$ is f-inductive starting at a, then the least fixed point of f above a is in P.*

Proof Here is the idea. Using the hypotheses it is not hard to show all members of the ordinal sequence constructed in the proof of Proposition 5.2 are in P, so then the least fixed point of f above a is in P, since it is in the ordinal sequence. □

We note without proof that the Corollary dualizes to greatest fixed points.

There is a variation on the Knaster-Tarski result that we will need. Instead of finding a fixed point for a function, one seeks *two* values between which it oscillates. This requires not monotonicity, but *anti-monotonicity*.

Definition 5.5 *Let $\langle L, \leq \rangle$ be a lattice. A mapping $f : L \to L$ is* anti-monotone *if it is order reversing, $x \leq y$ implies $f(y) \leq f(x)$. An* alternating fixpoint pair *for f is a pair x and y such that $f(x) = y$ and $f(y) = x$. An alternating fixpoint pair is* extremal *if any other alternating fixpoint pair is between them.*

Proposition 5.6 *Every anti-monotone mapping on a complete lattice has an extremal alternating fixpoint pair.*

Proof This can be proved directly, but there is a simple derivation from Proposition 5.2. Let $\langle L, \leq\rangle$ be a complete lattice and $f : L \to L$ be anti-monotone. The function $f^2(x) = f(f(x))$ is obviously monotone, and so has fixed points; call the set of them $\mathcal{F}$. Since f^2 has a least and a greatest fixed point then $\mathcal{F}$ has a least member, s, and a greatest, S. We show these are an extremal alternating fixpoint pair for f.

We first show that f maps $\mathcal{F}$ to $\mathcal{F}$. For, suppose $x \in \mathcal{F}$; then $f^2(f(x)) = f(f^2(x)) = f(x)$, so $f(x)$ is a fixed point of f^2, and hence in $\mathcal{F}$.

Since $s \in \mathcal{F}$ then $f(s) \in \mathcal{F}$. We show $f(s)$ is largest in $\mathcal{F}$, and hence is S. Let x be an arbitrary member of $\mathcal{F}$. Then $f(x) \in \mathcal{F}$, so $s \leq f(x)$. By anti-monotonicity, $f(f(x)) \leq f(s)$, so $x \leq f(s)$. Since x was arbitrary, $f(s) = S$. Then also $f(S) = f(f(s)) = s$, so f alternates between s and S.

Finally s and S are extremal. Suppose x and y are an alternating fixpoint pair for f. Then $f^2(x) = f(f(x)) = f(y) = x$ so x is a fixpoint of f^2. Similarly for y. Then x and y are in $\mathcal{F}$ and so are between s and S. □

We now move on to our main subject, bilattices. Recall we understand the term to include completeness with respect to each lattice order, as well as infinitary interlacing conditions connecting the orderings. We begin with a technical lemma whose proof is a prime example of interlacing conditions at work. It is essential for the subsequent results.

Lemma 5.7 *(From Avron (1996)) In a bilattice $\mathcal{B}$:*

1. *For any $a \leq_t b$ we have $a \leq_t x \leq_t b$ if and only if $a \otimes b \leq_k x \leq_k a \oplus b$;*
2. *For any $a \leq_k b$ we have $a \leq_k x \leq_k b$ if and only if $a \wedge b \leq_t x \leq_t a \vee b$.*

Proof We show the first item; the second is similar. Assume throughout that $a \leq_t b$.

1. *left-right*: Suppose $a \leq_t x \leq_t b$. By interlacing, $a \otimes (a \otimes b) \leq_t x \otimes (a \otimes b) \leq_t b \otimes (a \otimes b)$, or $a \otimes b \leq_t x \otimes (a \otimes b) \leq_t a \otimes b$. Then $x \otimes (a \otimes b) = a \otimes b$, and so $a \otimes b \leq_k x$. Similarly $x \leq_k a \oplus b$ by an argument using $\oplus$ instead of $\otimes$.
2. *right-left*: Suppose $a \otimes b \leq_k x \leq_k a \oplus b$. By interlacing, $a \wedge (a \otimes b) \leq_k a \wedge x \leq_k a \wedge (a \oplus b)$. Since $a \leq_t b$, by interlacing $a \otimes a \leq_t a \otimes b$, or $a \leq_t a \otimes b$, and hence $a \wedge (a \otimes b) = a$. Similarly $a \wedge (a \oplus b) = a$. Combining all this, $a = a \wedge (a \otimes b) \leq_k a \wedge x \leq_k a \wedge (a \oplus b) = a$. Then $a \wedge x = a$, and hence $a \leq_t x$. There is a similar proof that $x \leq_t b$. □

Proposition 5.8 *Suppose $\langle \mathcal{B}, \leq_t, \leq_k\rangle$ is a bilattice, and $f : \mathcal{B} \to \mathcal{B}$ is a mapping that is monotone in* both *orderings. By the Knaster-Tarski theorem f has a least and greatest fixed point with respect to each ordering. We use p_t, P_t for smallest and biggest fixed point with respect to $\leq_t$, and p_k, P_k for smallest and biggest with respect to $\leq_k$. The following hold: $P_t \otimes p_t = p_k$, $P_t \oplus p_t = P_k$, $P_k \wedge p_k = p_t$, $P_k \vee p_k = P_t$.*

Proof p_t and P_t are least and greatest fixed points with respect to $\leq_t$, and p_k is a fixed point, so $p_t \leq_t p_k \leq_t P_t$. By Lemma 5.7, $p_t \otimes P_t \leq_k p_k \leq_k p_t \oplus P_t$. Since p_k is least fixed point under $\leq_k$, and p_t and P_t are fixed points, $p_k \leq_k p_t$ and $p_k \leq_k P_t$, so $p_k \leq_k p_t \otimes P_t$. Combining, $P_t \otimes p_t = p_k$. The other three are similar. □

As a special case of Proposition 5.8, let f be the identity map on $\mathcal{B}$. Obviously this is monotone in both bilattice orderings, and everything is a fixed point, so the least and greatest fixed points under $\leq_t$ are **f** and **t**, while least and greatest under $\leq_k$ are $\bot$ and $\top$. Now Proposition 2.6 is an immediate consequence.

A version of Proposition 5.8 that we will need in Section 9.3 can now be obtained for the alternating case.

Proposition 5.9 *Suppose* $\langle \mathcal{B}, \leq_t, \leq_k \rangle$ *is a bilattice and* $f : \mathcal{B} \to \mathcal{B}$ *is a mapping that is monotone with respect to* $\leq_k$ *but anti-monotone with respect to* $\leq_t$*. Let* p_k *and* P_k *be the least and greatest fixed points with respect to* $\leq_k$*, and let* p_t *and* P_t *be the smaller and the larger of the extremal alternating fixpoint pair with respect to* $\leq_t$*. The following hold:* $P_t \otimes p_t = p_k$, $P_t \oplus p_t = P_k$. $P_k \wedge p_k = p_t$, $P_k \vee p_k = P_t$.

Proof Very simply, the mapping f^2 is monotone with respect to both orderings, and now use Proposition 5.8. □

6 Adding a Truth Predicate

We want to have a formal language that can speak about its own syntax. There are several ways of doing this; we use the traditional one, with an arithmetic language that talks about itself via Gödel numbering. Details are not particularly important. We take advantage of a point that Kripke has emphasized, and use relation symbols for addition and multiplication instead of function symbols. This has the nice consequence that every number has a unique term in the language that names it. In the following discussion we use symbols informally as names for themselves.

From now on assume $\mathcal{L}$ is a first-order language with a constant symbol 0 and a unary function symbol for successor—we write successor of t as t^+. The only terms of the language are the open ones, $x^{++\cdots+}$ where x is a variable, and the closed ones, $0^{++\cdots+}$, *numerals*. The number of successor symbols is allowed to be 0. We also assume we have a binary relation symbol =, and two ternary relation symbols A and M, intended to represent addition and multiplication.

First-order formulas are built up using propositional connectives, $\wedge$, $\vee$, $\neg$, and quantifiers, $\forall$, $\exists$, in the standard way and, if modality is being considered, also using $\Box$ and $\Diamond$. When evaluating formulas we will make use of the standard model of arithmetic to assign meaning to the non-logical symbols, and to supply the range for the quantifiers. It is possible to have additional symbols to represent 'real world' things or relations, but this adds complexity without providing any fundamental new insights, so we avoid it here. Truth values for sentences of $\mathcal{L}$ will be in a bilattice, whose choice can be quite arbitrary. Keep $\mathcal{FOUR}$ in mind as an example.

In addition to the machinery described above, $\mathcal{L}$ has a unary predicate symbol, T, intended to represent *is true*. Using it $\mathcal{L}$ can "talk about" the truth of its own sentences, using Gödel numbering. We ignore the details of Gödel numbering except that, for convenience, we assume every sentence (closed formula) has a Gödel number, and every number is the Gödel number of some sentence of $\mathcal{L}$. Numerals, 0, 0^{+}, 0^{++}, ..., provide a unique representation for all numbers and so, via Gödel numbering, of every sentence. Let X be a sentence of $\mathcal{L}$; by $\ulcorner X \urcorner$ we mean the numeral representing the Gödel number of X. If T is to act like a truth predicate we should have equivalence of X and $\mathsf{T}(\ulcorner X \urcorner)$ for every sentence X, that is, they should always have the same value in the bilattice that supplies our truth values. Since we have enough machinery to construct a liar sentence asserting its own non-truth, this can't be done if we use classical logic. Hence the historic move to three-valued logics, and here to bilattices.

Examples 4.7 and 4.9 looked at bilattices and first-order languages. In those examples we allowed members of domains to appear in sentences, as a convenient way of handling satisfiability. We no longer need to do this. We have numerals in $\mathcal{L}$, the natural numbers are the only domain we are interested in, and so we have names in $\mathcal{L}$ for the entire domain. Further, since arithmetic machinery will always be understood as if in the standard model for arithmetic, it is only the truth predicate T that can vary. So valuations will not assign values to all atomic sentences, but only to those involving T.

Definition 6.1 (T Valuation) $\mathcal{A}(\mathsf{T})$ *is the set of atomic sentences of the form* $\mathsf{T}(0^{++\cdots+})$ *in the language* $\mathcal{L}$*. For any bilattice* $\mathcal{B}$*, a* T valuation in $\mathcal{B}$ *is a member of the function space bilattice* $\mathcal{B}^{\mathcal{A}(\mathsf{T})}$ *(Definition 3.3),*

A T valuation in $\mathcal{B}$ simply assigns a truth value in $\mathcal{B}$ to every atomic sentence involving T. Example 4.9 illustrates how the definition above covers the modal case, since the bilattice to which T valuations map can be of the form $\mathcal{B}^{\mathcal{G}}$ where $\mathcal{G}$ is the set of possible worlds of a modal frame. We will assume that if a language $\mathcal{L}$ includes modal operators, then the underlying bilattices are of this form.

Examples 4.7 and 4.9 showed how valuations extend from the atomic level to *all* sentences, that is, to members of the function space bilattice $\mathcal{B}^{\mathcal{L}}$. But in Kripke (1975) *three different* extensions were discussed, these have generalizations here, and others are possible. The following allows us flexibility while placing minimal requirements on extensions. Specific examples come in Section 8.

Definition 6.2 (Valuation Extension) *A* valuation extension *in bilattice* $\mathcal{B}$ *is a mapping* $e : \mathcal{B}^{\mathcal{A}(\mathsf{T})} \to \mathcal{B}^{\mathcal{L}}$*, where* $\mathcal{L}$ *is understood to be the set of* sentences *of our formal language. For* $v \in \mathcal{B}^{\mathcal{A}(\mathsf{T})}$ *we write the image under* e *as* v^{e}.

Our goal is to show that for several intuitively acceptable valuation extensions e there are valuations $v \in \mathcal{B}^{\mathcal{A}(\mathsf{T})}$ such that for each sentence X, $v^{e}(X)$ and $v^{e}(\mathsf{T}(\ulcorner X \urcorner))$ are the same. If we can do this then we have ways in which T can behave like a truth predicate. Following ideas that started with Kripke (1975) and Martin and Woodruff (1975), we use what is often called a *truth revision operator* where, if v is a T valuation, then applying a truth revision operator to v produces another T valuation that explicitly incorporates some of the information that was implicit in v itself. We then look for valuations that are unchanged by this process.

Definition 6.3 (Truth Revision Operator) *Let e be a valuation extension in $\mathcal{B}$. An e* truth revision operator *on bilattice $\mathcal{B}$ is a mapping $\varphi^e_{\mathcal{B}} : \mathcal{B}^{\mathcal{A}(\mathsf{T})} \to \mathcal{B}^{\mathcal{A}(\mathsf{T})}$ such that for each* T *valuation v in $\mathcal{B}$, $\varphi^e_{\mathcal{B}}(v)$ is the* T *valuation in $\mathcal{B}$ such that $\varphi^e_{\mathcal{B}}(v)(\mathsf{T}(\ulcorner X \urcorner)) = v^e(X)$ for every sentence X of $\mathcal{L}$.*

A fixed point of truth revision operator $\varphi^e_{\mathcal{B}}$ is a T valuation v such that $v = \varphi^e_{\mathcal{B}}(v)$ and thus, for every sentence X, $v(\mathsf{T}(\ulcorner X \urcorner)) = \varphi^e_{\mathcal{B}}(v)(\mathsf{T}(\ulcorner X \urcorner)) = v^e(X)$. Then with respect to fixed point valuations, the truth predicate T does have the fundamental property we want a truth predicate to have. Of course liar sentences show this can't happen in a classical framework, but the additional machinery of bilattices gives us exactly what we need. To show fixed points exist in a bilattice, the key thing needed is monotonicity, and then we can make use of Knaster-Tarski. We saw results called monotonicity in Propositions 4.6, 4.8, and 4.10. The following is a general characterization, bringing the T predicate into things.

Definition 6.4 (Monotonicity Property) *We say the* monotonicity property *holds for valuation extension e in bilattice $\mathcal{B}$ provided, for any* T *valuations, if $v_1 \le_k v_2$ in $\mathcal{B}^{\mathcal{A}(\mathsf{T})}$ then $v_1^e \le_k v_2^e$ in $\mathcal{B}^{\mathcal{L}}$.*

Proposition 6.5 *If valuation extension e has the monotonicity property in bilattice $\mathcal{B}$ then the truth revision operator $\varphi^e_{\mathcal{B}}$ is monotonic in the information ordering in $\mathcal{B}^{\mathcal{A}(\mathsf{T})}$, that is, $v_1 \le_k v_2$ implies $\varphi^e_{\mathcal{B}}(v_1) \le_k \varphi^e_{\mathcal{B}}(v_2)$.*

Proof Assume e has the monotonicity property, and $v_1 \le_k v_2$. To show that $\varphi^e_{\mathcal{B}}(v_1) \le_k \varphi^e_{\mathcal{B}}(v_2)$ we show that for every sentence X we have $\varphi^e_{\mathcal{B}}(v_1)(\mathsf{T}(\ulcorner X \urcorner)) \le_k \varphi^e_{\mathcal{B}}(v_2)(\mathsf{T}(\ulcorner X \urcorner))$. But by Definition 6.3, this is equivalent to $v_1^e(X) \le_k v_2^e(X)$, and we have this by the monotonicity property, Definition 6.4, since we have $v_1 \le_k v_2$.□

Corollary 6.6 *In bilattice $\mathcal{B}$, if valuation extension e has the monotonicity property then the corresponding truth revision operator $\varphi^e_{\mathcal{B}}$ has fixed points, in particular it has smallest and greatest ones with respect to the $\le_k$ bilattice ordering.*

This Corollary is immediate by the Knaster-Tarski result, Proposition 5.2, though at the moment we don't officially know that there are valuation extensions having the monotonicity property. In fact there are several, as will be shown in Section 8.

7 Conflation and Consistency

When looking at $\mathcal{FOUR}$ in Figure 1 certain substructures naturally present themselves. The subset $\{\mathbf{f}, \mathbf{t}\}$ suggests classical logic; the truth operations $\wedge, \vee, \neg$ restricted to $\{\mathbf{f}, \mathbf{t}\}$ have exactly the classical behavior. Similarly the subset $\{\mathbf{f}, \bot, \mathbf{t}\}$ behaves like the strong Kleene logic K_3. Dually, the subset $\{\mathbf{f}, \top, \mathbf{t}\}$ can be thought of as like

Priest's logic of paradox LP. In some sense the lower values in $\mathcal{FOUR}$ represent consistency, while the upper values represent a kind of dual notion. (We will use "lower" and "upper" informally here, corresponding to position in the $\leq_k$ bilattice ordering.) As it turns out, analogs of these structures can be found in bilattices generally provided we have a conflation operator, something that has played little role for us up to this point.

As motivation let S be a set of agents (recall Example 4.1), with $\mathcal{P}(S)$ as the collection of all subsets, and with $\subseteq$ as a partial ordering. This is a De Morgan lattice with complementation, $X \mapsto (S - X)$, as involution (Definition 3.2). The bilattice product, $\mathcal{P}(S) \odot \mathcal{P}(S)$, is a bilattice with conflation. In it a value is a pair $\langle F, A\rangle$ where F is the set of agents with an opinion *for* some proposition, and A is the set with an opinion *against*. An agent may have an opinion both for and against, in which case $\langle F, A\rangle$ can be thought of as representing an inconsistency. An agent may have no opinion either way, in which case $\langle F, A\rangle$ is an incomplete bilattice value. Or it may be that every agent falls into exactly one of F or A, and so we have a kind of exactness of opinions. We incorporate these ideas into bilattices in a structural way. If every agent falls into exactly one of the *for* or *against* categories the corresponding bilattice value is of the form $\langle F, \overline{F}\rangle$, and values like these are their own conflations. If no agent is both *for* and *against*, the bilattice value is of the form $\langle F, A\rangle$ where $F \cap A = \emptyset$, or equivalently where $F \subseteq \overline{A}$, or $A \subseteq \overline{F}$. Using conflation this can be expressed quite simply as $\langle F, A\rangle \leq_k -\langle F, A\rangle$. That some agent is both *for* and *against* corresponds to $-\langle F, A\rangle \leq_k \langle F, A\rangle$ but $-\langle F, A\rangle \neq \langle F, A\rangle$. This leads us to the following.

Definition 7.1 (Consistent, Exact, Anticonsistent) *A member x of a bilattice with conflation is* consistent *if $x \leq_k -x$,* anticonsistent *if $-x \leq_k x$, and* exact *if $x = -x$.*

This three way division does not always exhaust a bilattice. $\mathcal{SIXTEEN}$, in Figure 3, has a conflation, $\langle\bot, \top\rangle$ and $\langle\top, \bot\rangle$ are conflations of each other but are not comparable in the $\leq_k$ ordering, and so are neither exact, consistent, nor anticonsistent. Such values merit further investigation, but not here. In $\mathcal{FOUR}$ the consistent members correspond to K_3, the anticonsistent members to LP, and the exact members to classical logic. The results that follow say these categories, exhaustive or not, always give us coherent subsystems of bilattices with conflation.

Proposition 7.2 *In a bilattice with conflation we have the following.*

1. *Each of the consistent, anticonsistent, and exact subsets are closed under truth operations $\wedge$, $\vee$, and $\neg$, as well as under infinitary operations $\bigwedge$ and $\bigvee$.*
2. *The consistent values are closed under chain sups with respect to the $\leq_k$ ordering, Definition 5.3.*
3. *Dually the anticonsistent values are closed under chain infs.*
4. *Anything below a consistent value, in the $\leq_k$ ordering, is consistent.*

Proof We work in a biattice with conflation.

1. If x and y are consistent, $x \leq_k -x$ and $y \leq_k -y$, then by interlacing, $x \wedge y \leq_k -x \wedge -y$, and by Proposition 2.5 $-x \wedge -y = -(x \wedge y)$, so $x \wedge y \leq_k -(x \wedge y)$, and

hence $x \wedge y$ is consistent. Using infinitary interlacing a similar proof gives that if every member of S is consistent then $\bigwedge S$ is consistent. The cases with $\vee$ and $\bigvee$ are similar. If $x \leq_k -x$ then $\neg x \leq_k \neg - x$, so $\neg x \leq_k -\neg x$ since negation and conflation commute, so $\neg x$ is consistent. Anticonsistent and exact are similar.

2. Suppose S is a chain of consistent bilattice values. We first show that for any $a, b \in S$ we have $a \leq_k -b$. Since both a and b are in the chain S, either $a \leq_k b$ or $b \leq_k a$. Since members of S are consistent, $a \leq_k -a$ and $b \leq -b$. If $a \leq_k b$, since $b \leq_k -b$, then $a \leq_k -b$. If $b \leq_k a$, since $a \leq_k -a$, then $b \leq_k -a$ so $a \leq_k -b$. Then either way, $a \leq_k -b$.
   Let $-S = \{-x \mid x \in S\}$. Now let a and b be arbitrary members of S. Then $-b$ is an arbitrary member of $-S$, and by what was just shown, $a \leq_k -b$. Since $-b$ is arbitrary, $a \leq -b$ for every $-b \in -S$, so a is a lower bound for $-S$, and hence $a \leq_k \prod -S$. Since a was arbitrary, $\prod -S$ is an upper bound for S, and so $\sum S \leq_k \prod -S$. Then by Proposition 2.5, $\sum S \leq_k -\sum S$, and so $\sum S$ is consistent.
3. By an argument dual to the previous one.
4. Suppose a is consistent (so $a \leq_k -a$) and $b \leq_k a$ (and so $-a \leq_k -b$). Then $b \leq_k a \leq_k -a \leq_k -b$, so b is consistent. □

Next, something that will play a role when we come to supervaluations in Section 8.4.

Proposition 7.3 *In a bilattice with conflation, if x is consistent then $x \leq_k y$ for some exact y, and $x = \prod\{y \mid y$ is exact and $x \leq_k y\}$.*

Proof A direct proof of the second part is possible but ornate, to say the least. In Section 3.1 we noted that the bilattice product construction was universal. Then if we prove the result for product bilattices $L \odot L$, where L is any non-distributive De Morgan lattice, we have the result generally.

Assume $\langle x_1, x_2\rangle \in L \odot L$, as in Definition 3.1, and is consistent, so $\langle x_1, x_2\rangle \leq_k -\langle x_1, x_2\rangle = \langle \overline{x_2}, \overline{x_1}\rangle$, which unwinds to $x_1 \leq \overline{x_2}$, or equivalently, $x_2 \leq \overline{x_1}$. Let $S = \{\langle y_1, y_2\rangle \mid \langle y_1, y_2\rangle$ is exact and $\langle x_1, x_2\rangle \leq_k \langle y_1, y_2\rangle\}$. We show $\langle x_1, x_2\rangle = \prod S$.

Trivially $\langle x_1, x_2\rangle \leq_k \prod S$. Also $\langle x_1, \overline{x_1}\rangle$ is easily seen to be exact and $\langle x_1, x_2\rangle \leq_k \langle x_1, \overline{x_1}\rangle$ because $x_2 \leq \overline{x_1}$, so $\langle x_1, \overline{x_1}\rangle \in S$ (incidentally giving the first part of the Proposition). Similarly $\langle \overline{x_2}, x_2\rangle \in S$. And $\langle x_1, \overline{x_1}\rangle \otimes \langle \overline{x_2}, x_2\rangle = \langle x_1 \sqcap \overline{x_2}, \overline{x_1} \sqcap x_2\rangle = \langle x_1, x_2\rangle$ because $x_1 \leq \overline{x_2}$. Then $\prod S \leq_k \langle x_1, \overline{x_1}\rangle \otimes \langle \overline{x_2}, x_2\rangle = \langle x_1, x_2\rangle$.

Part 1 of Proposition 7.2 says that bilattices with conflation have subsystems that are natural generalizations of K_3, LP, and classical logic. Kripke's work was entirely within the consistent part of $\mathcal{FOUR}$. We now discuss what is needed to recover his results—what conditions keep us within the consistent parts of bilattices. Note that if $\mathcal{B}$ is a bilattice with conflation, so is $\mathcal{B}^{\mathcal{A}(\mathsf{T})}$; and $v \in \mathcal{B}^{\mathcal{A}(\mathsf{T})}$ is consistent if and only if $v(\mathsf{T}(0^{++\cdots+}))$ is consistent in $\mathcal{B}$ for every numeral, and similarly for anticonsistency.

Definition 7.4 *Let $\mathcal{B}$ be a bilattice with conflation. Valuation extension e* preserves consistency *if, for every consistent $v \in \mathcal{B}^{\mathcal{A}(\mathsf{T})}$, the* T *valuation $\varphi^e_{\mathcal{B}}(v)$ is consistent, where $\varphi^e_{\mathcal{B}}$ is the truth revision operator. Similarly for* preserving anticonsistency.

Proposition 7.5 *Let $\mathcal{B}$ be a bilattice with conflation. If valuation extension e preserves consistency then the least fixed point of operator $\varphi^e_{\mathcal{B}}$ is consistent in $\mathcal{B}^{\mathcal{A}(\mathcal{T})}$. Similarly if e preserves anticonsistency then the greatest fixed point is anticonsistent.*

Proof We show the set of consistent members of the bilattice $\mathcal{B}^{\mathcal{A}(\mathsf{T})}$ is $\varphi^e_{\mathcal{B}}$ inductive starting at $\bot$ (using the $\leq_k$ ordering), Definition 5.3.

The least member $\bot$ of $\mathcal{B}^{\mathcal{A}(\mathsf{T})}$ is consistent because it is below everything, in particular its conflation. If v is a T valuation that is consistent in $\mathcal{B}^{\mathcal{A}(\mathsf{T})}$ then $\varphi^e_{\mathcal{B}}(v)$ is consistent because we are assuming preservation of consistency. Finally the consistent values are closed under chain sups by part 2 of Proposition 7.2.

The result follows by Corollary 5.4. And the entire argument dualizes. □

8 Particular Valuation Extensions

Valuation extensions were introduced in Section 6, specifically in Definition 6.2. Many different ones are plausible and interesting. We look at four of them.

8.1 Kleene's Strong Three Valued Logic, Generalized

We use s for the valuation extension map introduced below to suggest *strong*, as in Kleene's strong three valued logic, K_3. We emphasize the Kleene logic since it played a big role in Kripke (1975), but what we do also generalizes Graham Priest's three-valued Logic of Paradox, LP, and the Belnap-Dunn $\mathcal{FOUR}$. It might be helpful to review Example 4.9.

Definition 8.1 (Generalized Strong Kleene) *Let $\mathcal{B}$ be a bilattice. The strong Kleene valuation extension is the mapping $s : \mathcal{B}^{\mathcal{A}(\mathsf{T})} \to \mathcal{B}^{\mathcal{L}}$ where, for $v \in \mathcal{B}^{\mathcal{A}(\mathsf{T})}$, $v^s \in \mathcal{B}^{\mathcal{L}}$ simply uses the operations of $\mathcal{B}$ from the $\leq_t$ ordering, plus arithmetic facts.*

1. *For atomic sentences involving T, $v^s(\mathsf{T}(0^{++\cdots+})) = v(\mathsf{T}(0^{++\cdots+}))$.*
2. *If X is an atomic sentence not involving T, it will be a sentence of arithmetic. It is understood as in the standard model for arithmetic. Specifically, for a, b, and c being numerals that name numbers a°, b°, and c° in the standard model:*
   a. *$v^s(A(a,b,c)) = \mathbf{t}$ in $\mathcal{B}$ if $a^\circ + b^\circ = c^\circ$, and otherwise $v^s(A(a,b,c)) = \mathbf{f}$;*
   b. *$v^s(M(a,b,c)) = \mathbf{t}$ in $\mathcal{B}$ if $a^\circ \times b^\circ = c^\circ$, and otherwise $v^s(M(a,b,c)) = \mathbf{f}$;*
   c. *$v^s(a = b) = \mathbf{t}$ if $a^\circ = b^\circ$, and otherwise $v^s(a = b) = \mathbf{f}$.*
3. *$v^s(X \wedge Y) = v^s(X) \wedge v^s(Y)$*
4. *$v^s(X \vee Y) = v^s(X) \vee v^s(Y)$*
5. *$v^s(\neg X) = \neg v^s(X)$*
6. *$v^s(\forall x F(x)) = \bigwedge\{v^s(F(t)) \mid t \text{ is a numeral}\}$*
7. *$v^s(\exists x F(x)) = \bigvee\{v^s(F(t)) \mid t \text{ is a numeral}\}$*

If language $\mathcal{L}$ is modal, we add the following, for $\mathcal{F} = \langle \mathcal{G}, \mathcal{R} \rangle$ a modal frame.

8. $v^s(\Box X)(\Gamma) = \bigwedge\{v^s(X)(\Delta) \mid \Delta \in \mathcal{G} \text{ and } \Gamma\mathcal{R}\Delta\}$
9. $v^s(\Diamond X)(\Gamma) = \bigvee\{v^s(X)(\Delta) \mid \Delta \in \mathcal{G} \text{ and } \Gamma\mathcal{R}\Delta\}$

This gives us the monotonicity we need, with or without modal operators, Definition 6.4. Note that v and v^s agree on atomic sentences $\mathsf{T}(0^{++\cdots+})$.

Lemma 8.2 (Strong Kleene Monotonicity) *If $v_1 \leq_k v_2$ in $\mathcal{B}^{\mathcal{A}(\mathsf{T})}$, then $v_1^s \leq_k v_2^s$ in $\mathcal{B}^{\mathcal{L}}$.*

Proof Assume $v_1 \leq_k v_2$ and show $v_1^s(X) \leq_k v_2^s(X)$ by induction on the complexity of sentence X.

At the atomic level we have the result for sentences of the form $\mathsf{T}(0^{++\cdots+})$ because $v_1 \leq_k v_2$. We have it for arithmetic sentences because v_1^s and v_2^s agree on them.

For the induction step, if we have $X \wedge Y$ and the result is known for simpler formulas, then using the interlacing conditions, $v_1^s(X \wedge Y) = v_1^s(X) \wedge v_1^s(Y) \leq_k v_2^s(X) \wedge v_2^s(Y) = v_2^s(X \wedge Y)$. Similarly for $\vee$, while $\neg$ follows using the negation condition that $x \leq_k y$ implies $\neg x \leq_k \neg y$. The result follows in the quantifier cases using infinitary interlacing, and similarly for the modal cases. □

Now by Corollary 6.6, a strong Kleene truth revision operator has a smallest and a biggest fixed point in the bilattice $\mathcal{B}^{\mathcal{A}(\mathsf{T})}$ of T valuations in a bilattice $\mathcal{B}$. We can go further. The first part of what follows covers Kripke's use of Kleene's K_3, and the second part says how Priest's LP figures in.

Proposition 8.3 *For a bilattice $\mathcal{B}$ with conflation, the least fixed point of the strong Kleene truth revision operator is consistent and the greatest is anticonsistent.*

Proof Using Proposition 7.5, for the first part it is enough to show that $\varphi_{\mathcal{B}}^s$ preserves consistency, and for the second part that it preserves anticonsistency.

Suppose v is a T valuation that is consistent. To show $\varphi_{\mathcal{B}}^s(v)$ is consistent we show that, for every sentence X, $\varphi_{\mathcal{B}}^s(v)(\mathsf{T}(\ulcorner X \urcorner))$ is a consistent member of $\mathcal{B}$. Since v is consistent then $v(\mathsf{T}(t))$ is a consistent member of $\mathcal{B}$ for each numeral t. The value of v on atomic sentences of arithmetic must be either **t** or **f**, both of which are consistent. Then v^s maps all atomic sentences to consistent values in $\mathcal{B}$, so using Proposition 7.2 part 1, v^s maps every sentence of $\mathcal{L}$ to a consistent truth value. In particular, $v^s(X)$ is consistent, but this is $\varphi_{\mathcal{B}}^s(v)(\mathsf{T}(\ulcorner X \urcorner))$.

Anticonsistency is similar and is omitted. □

We now look at what the strong Kleene fixed points can be like for some of the bilattices from Section 4. Since arithmetic always behaves as it does in the standard model, we can use the Gödel fixed point theorem: for any formula $F(x)$ of language $\mathcal{L}$ having one free variable, a sentence X can be constructed such that $F(\ulcorner X \urcorner)$ and X are equivalent, in the sense that they always have the same value under v^s for any valuation v in a bilattice.

A *liar sentence* L is equivalent to $\neg\mathsf{T}(\ulcorner L \urcorner)$—a Gödel fixed point of $\neg\mathsf{T}(x)$. In any bilattice $\mathcal{B}$, for any valuation v that is a fixed point of truth revision operator

$\varphi^s_{\mathcal{B}}$, it must be the case that v^s assigns to L a bilattice value that is its own negation. For, suppose v were a fixed point valuation. Then $v^s(\mathsf{T}(\ulcorner L \urcorner)) = v(\mathsf{T}(\ulcorner L \urcorner)) = \varphi^s_{\mathcal{B}}(v)(\mathsf{T}(\ulcorner L \urcorner)) = v^s(L)$. But since L and $\neg\mathsf{T}(\ulcorner L \urcorner)$ are equivalent, $v^s(\neg\mathsf{T}(\ulcorner L \urcorner)) = v^s(L)$, so $v^s(L) = \neg v^s(\mathsf{T}(\ulcorner L \urcorner))$. Then $v^s(L) = \neg v^s(L)$. It follows that in no bilattice $\mathcal{B}$ can truth revision operator $\varphi^s_{\mathcal{B}}$ have a fixed point valuation in which the liar sentence L takes on either value **f** or **t**, for neither of these are negations of themselves.

In $\mathcal{FOUR}$, by the Knaster-Tarski theorem the strong Kleene truth revision operator has a smallest and a greatest fixed point under $\leq_k$. The liar sentence L cannot be assigned either **t** or **f** by fixed points, so the values assigned to the liar in $\mathcal{FOUR}$ must be $\bot$ or $\top$. The smallest fixed point is consistent and the biggest is anticonsistent.

More complicated bilattices than $\mathcal{FOUR}$ have technical interest and may have philosophical interest as well. Consider Figure 2 showing the bilattice $\mathcal{NINE}$. In this the smallest and biggest fixed point valuations still give the liar the values $\bot = \langle 0, 0\rangle$ and $\top = \langle 1, 1\rangle$, as they did in $\mathcal{FOUR}$. But in $\mathcal{NINE}$ there is one more value that is its own negation, $\langle \frac{1}{2}, \frac{1}{2}\rangle$, which suggests it is a value that the liar could take on in some fixpoint. In fact this happens, and thus in $\mathcal{NINE}$ there are at least three fixpoint valuations that handle the liar. To show there actually is such a fixed point we can use the Knaster-Tarski Theorem 5.2. Here is the construction.

Suppose L and $\neg\mathsf{T}(\ulcorner L \urcorner)$ are equivalent. Let v be the valuation in $\mathcal{NINE}$ such that $v(\mathsf{T}(\ulcorner L \urcorner)) = \langle \frac{1}{2}, \frac{1}{2}\rangle$, and for every sentence X other than L, $v(\mathsf{T}(\ulcorner X \urcorner)) = \langle 0, 0\rangle$, that is, the value is $\bot$ of $\mathcal{NINE}$. To keep notation uncluttered, let $\varphi = \varphi^s_{\mathcal{NINE}}$ be the strong Kleene truth revision operator for $\mathcal{NINE}$. Then $\varphi(v)(\mathsf{T}(\ulcorner L \urcorner)) = v^s(L) = v^s(\neg\mathsf{T}(\ulcorner L \urcorner)) = \neg v^s(\mathsf{T}(\ulcorner L \urcorner)) = \neg\langle \frac{1}{2}, \frac{1}{2}\rangle = \langle \frac{1}{2}, \frac{1}{2}\rangle = v(\mathsf{T}(\ulcorner L \urcorner))$, so $v(\mathsf{T}(\ulcorner L \urcorner)) = \varphi(v)(\mathsf{T}(\ulcorner L \urcorner))$. On all other sentences X, $v(\mathsf{T}(\ulcorner X \urcorner)) \leq_k \varphi(v)(\mathsf{T}(\ulcorner X \urcorner))$ since $v(\mathsf{T}(\ulcorner X \urcorner)) = \bot$. Then $v \leq_k \varphi(v)$. Now by Knaster-Tarski 5.2, part 2, there is a fixed point valuation v_f of φ above v in the $\leq_k$ ordering. It is easy to check that every valuation in the Ordinal Sequence approximating to v_f will assign $\langle \frac{1}{2}, \frac{1}{2}\rangle$ to $\ulcorner L \urcorner$, hence this is what v_f assigns.

In Example 4.4 we briefly discussed a Fuzzy Bilattice, the bilattice product of the closed unit interval with itself. Taking this as the space of truth values gives us a kind of 'degree of confidence' structure. The truth revision operator in the bilattice of valuations in this structure has fixpoints. The construction just presented for $\mathcal{NINE}$ can easily be transferred. Every fuzzy truth value of the form $\langle a, a\rangle$ is a value the liar sentence can take on in some strong Kleene fixpoint using the fuzzy bilattice.

8.2 Kleene's Weak Three Valued Logic, Generalized

A three valued logic that originated in Bochvar (1938) was introduced independently in Kleene (1950) with a different motivation. Commonly called *Kleene's weak three valued logic*, it was one of the logics employed in Kripke (1975), and the only logic used in Martin and Woodruff (1975). It behaves classically on classical truth values, but simply assigns $\bot$ to any formula under a valuation whenever any atomic subformula is assigned $\bot$ by that valuation. It can be reformulated in a way that has an intuition extending beyond the three valued setting. This was foreshadowed in Fitting (1994), and given its present form in Fitting (2006) under the name "cut down operations". For recent work on cut down operations, see Ferguson (2015).

As motivation we turn once again to bilattices arising from the opinions of sets of agents, continuing the discussions from Section 7 and Example 4.1. Recall that the setting is a product bilattice $\mathcal{P}(S) \odot \mathcal{P}(S)$, where $\mathcal{P}$ is the power set operation, S is a set of agents, and the ordering on $\mathcal{P}(S)$ is $\subseteq$. In this bilattice a truth value is $\langle F, A\rangle$ with the F values informally thought of as those agents *for* and the A values as those *against* some proposition. But now in evaluating, say, $X \wedge Y$ suppose we only want to consider the opinions of agents on Y who have actually expressed an opinion on X, and similarly the other way around. That is, we want to "cut down" the sets of agents to the completely active ones. Suppose the truth value assigned to X is $\langle F, A\rangle$; then the active agents on X are $F \cup A$. The consensus of $\langle F \cup A, F \cup A\rangle$ with the truth value assigned to Y would give us the kind of restriction we need. But also, $\langle F \cup A, F \cup A\rangle = \langle F, A\rangle \oplus \langle A, F\rangle = \langle F, A\rangle \oplus \neg\langle F, A\rangle$. This leads us to the following, where the superscript w is meant to suggest *weak*, as in Kleene's *weak* logic.

Definition 8.4 (Cut Down Operations) *Let $\mathcal{B}$ be a bilattice with x and y as members, and S as a set of members.*

1. $\|x\| = x \oplus \neg x$
2. $x \wedge^w y = (x \wedge y) \otimes \|x\| \otimes \|y\|$
3. $x \vee^w y = (x \vee y) \otimes \|x\| \otimes \|y\|$
4. $\bigwedge^w S = (\bigwedge S) \otimes \prod\{\|x\| \mid x \in S\}$
5. $\bigvee^w S = (\bigvee S) \otimes \prod\{\|x\| \mid x \in S\}$

We did not define a weak negation. If we had, it would have been $\neg^w x = \neg x \otimes \|x\|$, but this is easily seen to be the same as $\neg x$.

The cut down operations restricted to the consistent part of $\mathcal{FOUR}$, $\{\mathbf{f}, \bot, \mathbf{t}\}$, are exactly the weak Kleene or Bochvar operations. Indeed, the motivation for the cut down operations suggests this should be so since this is the one agent case, and if the agent has no opinion about something, that lack of opinion propagates through any application of the weak operations. Of course our logic operations are actually on the whole of $\mathcal{FOUR}$, and the extension follows the familiar pattern. Throughout $\mathcal{FOUR}$, $\wedge^w$ and $\vee^w$ behave the same as $\wedge$ and $\vee$, except that any conjunction or disjunction with a component evaluating to $\bot$ itself evaluates to $\bot$.

Proposition 8.5 *In any bilattice with conflation the set of consistent values is closed under the truth operations $\wedge^w$, $\vee^w$, and under the infinitary operations $\bigwedge^w$ and $\bigvee^w$.*

Proof Suppose both x and y are consistent. By Proposition 7.1, $x \wedge y$ is consistent and so $x \wedge y \leq_k -(x \wedge y)$. Then,

$$\begin{aligned} x \wedge^w y &= (x \wedge y) \otimes \|x\| \otimes \|y\| \\ &\leq_k (x \wedge y) \leq_k -(x \wedge y) \\ &\leq_k -(x \wedge y) \oplus -\|x\| \oplus -\|y\| \\ &= -[(x \wedge y) \otimes \|x\| \otimes \|y\|] \\ &= -(x \wedge^w y) \end{aligned}$$

so $x \wedge^w y$ is consistent. The other cases are similar. ◻

Note that, unlike with our strong Kleene generalization in Proposition 7.2, we did not claim closure of the anticonsistent or exact part under our weak logic connectives.

Example 8.6 In the bilattice $\mathcal{SIXTEEN}$ from Figure 3 neither the anticonsistent nor the exact values are closed under $\wedge^w$. Recall, this is the bilattice product of $\mathcal{FOUR}$ from Figure 1 with itself, using the $\leq_t$ lattice ordering for $\mathcal{FOUR}$. With this ordering, $\mathcal{FOUR}$ is a De Morgan lattice with $\overline{\top} = \top$, $\overline{\bot} = \bot$, $\overline{\mathbf{t}} = \mathbf{f}$, and $\overline{\mathbf{f}} = \mathbf{t}$.

In $\mathcal{SIXTEEN}$ $-\langle \bot, \bot \rangle = \langle \overline{\bot}, \overline{\bot} \rangle = \langle \bot, \bot \rangle$ and similarly $-\langle \top, \top \rangle = \langle \top, \top \rangle$, so both $\langle \bot, \bot \rangle$ and $\langle \top, \top \rangle$ are exact, and hence are both anticonsistent and consistent as well. Now $\langle \bot, \bot \rangle \wedge^w \langle \top, \top \rangle = (\langle \bot, \bot \rangle \wedge \langle \top, \top \rangle) \otimes \|\langle \bot, \bot \rangle\| \otimes \|\langle \top, \top \rangle\| = \langle \mathbf{f}, \mathbf{t} \rangle \otimes \langle \bot, \bot \rangle \otimes \langle \top, \top \rangle = \langle \mathbf{f}, \mathbf{f} \rangle$. But $-\langle \mathbf{f}, \mathbf{f} \rangle = \langle \overline{\mathbf{f}}, \overline{\mathbf{f}} \rangle = \langle \mathbf{t}, \mathbf{t} \rangle$ so $\langle \mathbf{f}, \mathbf{f} \rangle$ is not exact, and it is not anticonsistent since $\langle \mathbf{t}, \mathbf{t} \rangle \not\leq_k \langle \mathbf{f}, \mathbf{f} \rangle$. In accord with Proposition 8.5, it should be consistent and, in fact, we do have $\langle \mathbf{f}, \mathbf{f} \rangle \leq_k \langle \mathbf{t}, \mathbf{t} \rangle$.

More generally, suppose $L \odot L$ is a bilattice product, where L is a non-distributive De Morgan lattice. It can be shown that, in $L \odot L$, the exact values are closed under the weak logic connectives $\wedge^w$ and $\vee^w$ if and only if, in L, $x \vee \overline{x} = 1$ and $x \wedge \overline{x} = 0$ for all x. This also allows us to conclude that the exact values of $\mathcal{SIXTEEN}$ are not closed under $\wedge^w$ and $\vee^w$.

We now define the obvious weak Kleene valuation extension, an instance of Definition 6.2. It is well behaved on the consistent part of a bilattice with conflation and at least makes sense outside that range, but its behavior there has not been investigated. The definition has been dualized in Szmuc (2018), producing what are called *track down* operations. These are well behaved on the anticonsistent part, and so produce a weak generalization of Priest's LP.

Definition 8.7 (Generalized Weak Kleene) *Let $\mathcal{B}$ be a bilattice. The weak Kleene valuation extension is the map $w : \mathcal{B}^{\mathcal{A}(\mathsf{T})} \to \mathcal{B}^{\mathcal{L}}$ where, for $v \in \mathcal{B}^{\mathcal{A}(\mathsf{T})}$, $v^w \in \mathcal{B}^{\mathcal{L}}$ uses the cut down operations on $\mathcal{B}$.*

1. *For atomic sentences involving* T, $v^w(\mathsf{T}(0^{++\cdots+})) = v^s(\mathsf{T}(0^{++\cdots+})) = v(\mathsf{T}(0^{++\cdots+}))$.
2. *For numerals a, b, and c:*
   a. $v^w(A(a, b, c)) = v^s(A(a, b, c))$;
   b. $v^w(M(a, b, c)) = v^s(M(a, b, c))$;
   c. $v^w(a = b) = v^s(a = b)$.
3. $v^w(X \wedge Y) = v^w(X) \wedge^w v^w(Y)$
4. $v^w(X \vee Y) = v^w(X) \vee^w v^w(Y)$

5. $v^w(\neg X) = \neg v^w(X)$
6. $v^w(\forall x F(x)) = \bigwedge^w \{v^w(F(t)) \mid t \text{ is a numeral}\}$
7. $v^w(\exists x F(x)) = \bigvee^w \{v^w(F(t)) \mid t \text{ is a numeral}\}$

If language $\mathcal{L}$ is modal, we add the following, where $\mathcal{F} = \langle \mathcal{G}, \mathcal{R} \rangle$ is a modal frame.

8. $v^w(\Box X)(\Gamma) = \bigwedge^w \{v^w(X)(\Delta) \mid \Delta \in \mathcal{G} \text{ and } \Gamma \mathcal{R} \Delta\}$
9. $v^w(\Diamond X)(\Gamma) = \bigvee^w \{v^w(X)(\Delta) \mid \Delta \in \mathcal{G} \text{ and } \Gamma \mathcal{R} \Delta\}$

As with strong Kleene, v and v^w agree on atomic sentences $\mathsf{T}(0^{++\cdots+})$. We now have the Monotonicity Property, Definition 6.4 by the following.

Lemma 8.8 (Weak Kleene Monotonicity) *If $v_1 \leq_k v_2$ in $\mathcal{B}^{\mathcal{A}(\mathsf{T})}$, then $v_1^w \leq_k v_2^w$ in $\mathcal{B}^{\mathcal{L}}$.*

Proof Assume $v_1 \leq_k v_2$. We show $v_1^w(X) \leq_k v_2^w(X)$ by induction on the complexity of sentence X. We only sketch one case and leave the rest to the reader. Suppose $v_1^w(X) \leq_k v_2^w(X)$ and $v_1^w(Y) \leq_k v_2^w(Y)$. We verify that $v_1^w(X \wedge Y) \leq_k v_2^w(X \wedge Y)$. In the following the inequality step is by the induction hypothesis, the lattice properties of $\oplus$ and $\otimes$, the behavior of $\neg$ with respect to $\leq_k$, and the interlacing conditions.

$$\begin{aligned} v_1^w(X \wedge Y) &= v_1^w(X) \wedge^w v_1^w(Y) \\ &= [v_1^w(X) \wedge v_1^w(Y)] \otimes \|v_1^w(X)\| \otimes \|v_1^w(Y)\| \\ &= [v_1^w(X) \wedge v_1^w(Y)] \otimes [v_1^w(X) \oplus \neg v_1^w(X)] \otimes [v_1^w(Y) \oplus \neg v_1^w(Y)] \\ &\leq_k [v_2^w(X) \wedge v_2^w(Y)] \otimes [v_2^w(X) \oplus \neg v_2^w(X)] \otimes [v_2^w(Y) \oplus \neg v_2^w(Y)] \\ &= v_2^w(X) \wedge v_2^w(Y) \end{aligned}$$

□

Now we have a weak Kleene truth revision operator $\varphi_{\mathcal{B}}^w$, from Definition 6.3. By Proposition 6.5, $\varphi_{\mathcal{B}}^w$ is monotonic, and hence has least and greatest fixed points. Using Propositions 8.5 and 8.3, the least fixed point of $\varphi_{\mathcal{B}}^w$ is consistent in a bilattice with conflation. Unlike with strong Kleene, however, the anticonsistent members of a bilattice are not always closed under the extended weak Kleene operations, so we have no general conclusions concerning the greatest fixed point.

Track down operations were mentioned earlier in this section, Szmuc (2018). Since their general behavior is dual to cut down operations, it is likely that they will yield a version of a fixed point truth theory generalizing one based on a weak version of Priest's LP. This does not seem to have been investigated yet.

8.3 Asymmetric Logics

There is a natural propositional three valued asymmetric logic that partially generalizes to bilattices. It seems to have had minimal discussion in the literature on theories of truth, though Kripke and others have talked about it informally, and tableau rules were considered in Fitting (1994). It is, in fact, common in programming languages, where it is generally connected with *lazy evaluation*. Suppose we have $X \wedge Y$, where the values of X and Y are to be determined by calls on procedures which may or may not terminate, and so may or may not return a value. Lazy evaluation calls on the X routine first and if that returns **f** then the Y routine is not called, but instead the overall program concludes the conjunction is false. If the X routine returns **t**, only then is a further call on the Y routine made. Either the X or the Y routine may not terminate, but because lazy evaluation works from left to right there is an asymmetric difference between the two. If X does not terminate there is no final value for the conjunction, but if X returns **f**, even if the Y routine does not terminate, the conjunction has value **f**. It is a logic in which commutativity does not hold.

It is possible to define natural generalizations of propositional connectives for an asymmetric logic using the cut down ideas from Section 8.2. Recall our earlier discussion involving a group of agents. In these terms, for the conjunction $X \wedge Y$, asymmetry would have us only considering the opinions of experts on Y who have said something about X, but not the other way around. This idea gives us the following, where the a is for *asymmetric*. (Again negation does not need any special treatment.)

Definition 8.9 (Asymmetric Operations) *Let $x, y \in \mathcal{B}$ where $\mathcal{B}$ is a bilattice. Using Definition 8.4 part 1:*

1. $x \wedge^a y = (x \wedge y) \otimes \|x\|$
2. $x \vee^a y = (x \vee y) \otimes \|x\|$

Figure 4 gives tables for $\wedge^a$ and $\vee^a$ when the bilattice is $\mathcal{FOUR}$. Curiously, restricted to the consistent values $\{\mathbf{f}, \bot, \mathbf{t}\}$ the distributive laws hold, something that programmers make use of all the time. But with all four values present distributivity fails. For instance, $A \wedge^a (B \vee^a C)$ and $(A \wedge^a C) \vee^a (B \wedge^a C)$ differ if A has value $\top$, B has value $\bot$, and C is either $\top$ or **t**.

$\wedge^a$	$\top$	**t**	**f**	$\bot$
$\top$	$\top$	$\top$	**f**	**f**
t	$\top$	**t**	**f**	$\bot$
f	**f**	**f**	**f**	**f**
$\bot$	$\bot$	$\bot$	$\bot$	$\bot$

$\vee^a$	$\top$	**t**	**f**	$\bot$
$\top$	$\top$	**t**	$\top$	**t**
t	**t**	**t**	**t**	**t**
f	$\top$	**t**	**f**	$\bot$
$\bot$	$\bot$	$\bot$	$\bot$	$\bot$

Fig. 4: Four-valued Tables for $\wedge^a$ and $\vee^a$

Quantifiers are a problem. Since our domains are, by convention, the natural numbers and we want a left/right bias, we might think of $\forall x F(x)$ as if it were $F(0) \wedge^a F(1) \wedge^a F(2) \wedge^a \ldots$. But numbers code sentences, so the simple $0, 1, 2, \ldots$ order of evaluation might not be what we want at all. To allow full flexibility we could think of quantification as infinite conjunction and disjunction with an order specified. Formally, remove quantifiers from $\mathcal{L}$ and introduce infinitary connectives, $\bigwedge_{i \in \omega} A_i$ and $\bigvee_{i \in \omega} A_i$, where the indexing specifies an ordering. We could then think of the infinitary conjunction as being interpreted asymmetrically as $A_1 \wedge^a (A_2 \otimes \|A_1\|) \wedge^a (A_3 \otimes \|A_1\| \otimes \|A_2\|) \wedge^a \ldots$. Similarly for disjunction. More formally, for a sequence $x_1, x_2, x_3, \ldots \in \mathcal{B}$ where $\mathcal{B}$ is a bilattice, we could introduce the infinitary operations $\bigwedge^a_{i\in\omega} x_i = \bigwedge_{i\in\omega}\{x_i \otimes \prod\{\|x_j\| \mid j < i\}\}$ and $\bigvee^a_{i\in\omega} x_i = \bigvee_{i\in\omega}\{x_i \otimes \prod\{\|x_j\| \mid j < i\}\}$. (If $i = 1$, $\prod \emptyset$ occurs. This is $\top$, and $x_1 \otimes \top = x_1$.)

The asymmetric propositional operations from Definition 8.9 combined with the infinitary asymmetric replacements for quantification as just described work well up to a point. It is easy to show we have monotonicity results similar to Lemmas 8.2 and 8.8. But we can *not* define a monotone truth revision operator, Definition 6.3, for the elementary reason that infinitary conjunctions and disjunctions do not have Gödel numbers, and so a truth revision definition can't get off the ground.

Finite asymmetric conjunctions and disjunctions have some applications in natural languages. One might plausibly add them to the strong or weak Kleene systems, getting more expressive languages still having the same general fixed point behavior. One might allow infinite conjunctions whose conjuncts are given by a recursive function. Gödel numbering could perhaps then be reintroduced. We leave further investigation of this to others.

8.4 Supervaluations, Generalized

The least familiar three-valued logic used in Kripke (1975) is based on *supervaluations*, introduced in Van Fraassen (1966). Neither of Kleene's weak or strong three valued logics validate any classical tautologies since, for both logics, if all atomic formulas are assigned $\perp$ so is every formula. Supervaluations are specially crafted to validate tautologies, but the price paid is that supervaluations are not truth functional. We do not calculate via truth tables, but rather by a more complex method.

Kleene's strong logic, K_3, is three valued, but it extends naturally to all of the Belnap-Dunn four valued logic $\mathcal{FOUR}$. Moreover, every bilattice with conflation has a consistent substructure that generalizes K_3, and this always is a natual part of the entire of the bilattice. Kleene's weak three valued logic has generalizations that are most well-behaved on the consistent parts of bilattices with conflation, but the generalizations make sense for the entirety of bilattices. But supervaluation, in its original sense, is inherently a three valued logic. It lives on the consistent part of $\mathcal{FOUR}$, and makes no sense at all on the whole of it. There is a natural generalization of supervaluation to an arbitrary bilattice with conflation, but again only to the consistent part. Since we would like to have a complete lattice so that the Knaster-Tarski theorem can be applied, we will artificially extend our supervaluation generalization beyond the consistent part of a bilattice, to the entire structure. This does simplify the mathematics, but we are actually interested in behavior only on consistent values.

Here is the original version of supervaluations from Van Fraassen (1966), as used by Kripke. Suppose we have a K_3 valuation v mapping atomic sentences to $\{\mathbf{f}, \perp, \mathbf{t}\}$. That valuation can be 'raised' to a purely classical one by changing it on every atom that maps to $\bot$ so that instead the valuation maps the atom to one of the classical values, **f** or **t**. Call such a valuation a *totally defined extension* of v. For any given sentence X, the *supervaluation value* assigned to X by v is **t** if every totally defined extension of v maps X to **t**, the value is **f** if every totally defined extension of v maps X to **f**, and the value is $\bot$ otherwise, that is, when totally defined extensions disagree on X. To use bilattice terminology, we are working with valuations in the consistent part of $\mathcal{FOUR}$, and totally defined extensions map just to the exact part of $\mathcal{FOUR}$. Notice that if v is consistent in the bilattice of valuations in $\mathcal{FOUR}$ and w is a totally defined extension of v then w is exact and $v \leq_k w$. And the idea of valuations agreeing or disagreeing can be seen as an application of the bilattice consensus operation, $\otimes$. This should help motivate our generalization, which originated in Fitting (2006). Recall that outside the consistent part we will simply assign an arbitrary value, with $\top$ being the appropriate choice.

Definition 8.10 (Generalized Supervaluations) *Let $\mathcal{B}$ be a bilattice. The* supervaluation extension *is the map $sv : \mathcal{B}^{\mathcal{A}(\mathsf{T})} \to \mathcal{B}^{\mathcal{L}}$ where, for $v \in \mathcal{B}^{\mathcal{A}(\mathsf{T})}$, $v^{sv} \in \mathcal{B}^{\mathcal{L}}$ (using* sv *for* supervaluation*) is defined as follows.*

$$v^{sv} = \begin{cases} \prod\{w^s \mid w \in \mathcal{B}^{\mathcal{A}(\mathsf{T})} \text{ is exact and } v \leq_k w\} & \text{if } v \text{ is consistent} \\ \top & \text{if } v \text{ is not consistent} \end{cases}$$

In the non-consistent case, $\top$ is that of $\mathcal{B}^{\mathcal{L}}$, the function on $\mathcal{L}$ that is identically the $\top$ element of $\mathcal{B}$. In the consistent case this definition builds on top of the strong Kleene valuation extension. Recall that v is consistent in $\mathcal{B}^{\mathcal{A}(\mathsf{T})}$ if $v(\mathsf{T}(0^{++\cdots+}))$ is consistent in $\mathcal{B}$ for every numeral $0^{++\cdots+}$. In the consistent case the infinitary consensus operation is applied to a non-empty set by Proposition 7.3. We are not really interested in valuations that are not consistent—this case is here for technical reasons. But in the consistent case the output value, v^{sv}, will be consistent because, more generally, if S is any non-empty set of exact bilattice members, $\prod S$ is consistent. Very simply, S is non-empty, so let $s \in S$. Since s is exact it is consistent and since $\prod S \leq_k s$, $\prod S$ is consistent by Proposition 7.2 part 4. Also, in the consistent case v^{sv} assigns to arithmetic sentences the values they have in the standard model, because all strong Kleene extensions do this, and we have Proposition 7.3.

Lemma 8.11 (Supervaluation Monotonicity) *If $v_1 \leq_k v_2$ in $\mathcal{B}^{\mathcal{A}(\mathsf{T})}$, then $v_1^{sv} \leq_k v_2^{sv}$ in $\mathcal{B}^{\mathcal{L}}$.*

Proof Assume $v_1 \leq_k v_2$. If v_2 is not consistent, $v_2^{sv}(X) = \top$ for every X, so trivially $v_1^{sv}(X) \leq_k v_2^{sv}(X)$ for all sentences X. If v_2 is consistent so is v_1, by Proposition 7.2 part 4, so the first case of Definition 8.10 is used for both v_1^{sv} and v_2^{sv}. If w is an exact valuation and $v_2 \leq_k w$, then also $v_1 \leq_k w$, so

$$\{w \in \mathcal{B}^{\mathcal{A}(\mathsf{T})} \mid w \text{ is exact and } v_2 \leq_k w\} \subseteq \{w \in \mathcal{B}^{\mathcal{A}(\mathsf{T})} \mid w \text{ is exact and } v_1 \leq_k w\}$$

and so

$$
\begin{aligned}
v_1^{sv}(X) &= \prod\{w^s(X) \mid w \in \mathcal{B}^{\mathcal{A}(\mathsf{T})} \text{ and } w \text{ is exact and } v_1 \leq_k w\}\\
&\leq_k \prod\{w^s(X) \mid w \in \mathcal{B}^{\mathcal{A}(\mathsf{T})} \text{ and } w \text{ is exact and } v_2 \leq_k w\}\\
&= v_2^{sv}(X).
\end{aligned}
$$

□

The supervaluation truth revision operator using bilattice $\mathcal{B}$ (Definition 6.3) is denoted $\varphi_{\mathcal{B}}^{sv}$. It is monotone by Proposition 6.5, and has a smallest and a greatest fixed point under the $\leq_k$ ordering, by Corollary 6.6. The greatest fixed point of $\varphi_{\mathcal{B}}^{sv}$ is artificial, but the least fixed is consistent in a bilattice with conflation. In the consistent part of $\mathcal{FOUR}$ this treatment coincides with supervaluations as in Kripke, but the version here applies more generally, including to modal examples.

Supervaluations, in their original three-valued version, map every tautology to **t**. Things are not so simple for our generalization. Use the bilattice $\mathcal{NINE}$ from Example 4.2, let P be atomic, and consider the tautology $P \vee \neg P$ under the valuation v that maps P to $\langle 0, \frac{1}{2}\rangle$. The exact values above $\langle 0, \frac{1}{2}\rangle$ are $\langle 0, 1\rangle$ and $\langle \frac{1}{2}, \frac{1}{2}\rangle$. Let w_1 be an exact valuation mapping P to the first of these, and w_2 be one mapping P to the second. Then $w_1^s(P \vee \neg P) = \langle 0, 1\rangle \vee \neg\langle 0, 1\rangle = \langle 0, 1\rangle \vee \langle 1, 0\rangle = \langle 1, 0\rangle$ and $w_2^s(P \vee \neg P) = \langle \frac{1}{2}, \frac{1}{2}\rangle \vee \neg\langle \frac{1}{2}, \frac{1}{2}\rangle = \langle \frac{1}{2}, \frac{1}{2}\rangle$. But then $(w_1^s \otimes w_2^s)(P \vee \neg P) = \langle 1, 0\rangle \otimes \langle \frac{1}{2}, \frac{1}{2}\rangle = \langle \frac{1}{2}, 0\rangle$, which is not even exact. Further study of generalized supervaluations may be warranted.

As pointed out by a referee, there is a supervaluation dual called subvaluation—see Cobreros (2013) and Teijeiro (2020 forthcoming). It has been discussed for its applications to vagueness, which also supplied early motivation for supervaluations. It is likely that, when examined in the bilattice context, it will have behavior dual to that of supervaluations. So far there seems to have been no work that has been done on this.

9 Specialized Fixed Points

We have concentrated on least and greatest fixed points of truth revision operators, but there are other interesting varieties. Kripke (1975) introduced *intrinsic* fixed points, which also independently appeared in Manna and Shamir (1977) in a computer science context. *Maximal* fixed points were central in Martin and Woodruff (1975). *Alternating* fixed points come from Yablo (1985). There are probably others out there but these, at least, generalize to the bilattice family. Both Kripke (1975) and Martin and Woodruff (1975) worked with three valued logics. The counterpart to that in our approach is the *consistent* part of a bilattice, and a consistency restriction is noted in several parts of this section.

Terminology Convention 9.1 To keep clutter down we announce that throughout this entire section $\mathcal{B}$ is a bilattice with conflation. Also, throughout, we always assume that e is a valuation extension that has the monotonicity property (Definition 6.4) and preserves consistency (Definition 7.4). This includes generalized strong

Kleene, generalized weak Kleene, and generalized supervaluations. We assume $\varphi^e_{\mathcal{B}}$ is the corresponding truth revision operator. Finally $\mathcal{L}$ is our first-order arithmetic language, extended with T, and which may or may not contain modal operators.

9.1 Consistent Maximal Fixed Points

What we call consistent maximal fixed points are interesting in their own right, and play an important role with respect to intrinsic fixed points, which are discussed in Section 9.2.

Definition 9.2 (Consistent Maximal Fixed Point) *A fixed point v of $\varphi^e_{\mathcal{B}}$ is* consistent maximal *if it is consistent, and there is no other consistent fixed point w with $v \leq_k w$.*

The main thing we need is that there actually are consistent maximal fixed points. We recall Definition 7.4.

Proposition 9.3 *If v is a consistent T valuation for which $v \leq_k \varphi^e_{\mathcal{B}}(v)$ (a post-fixed point), then there is a consistent maximal fixed point of $\varphi^e_{\mathcal{B}}$ above v, and so consistent maximal fixed points exist.*

Proof Let $\mathcal{F} = \{w \in \mathcal{B}^{\mathcal{A}(\mathsf{T})} \mid w \text{ is consistent and } w \leq_k \varphi^e_{\mathcal{B}}(w)\}$. We show every member of $\mathcal{F}$ is below a consistent maximal fixed point, giving us the first part of the Proposition. The second part follows since $\bot$ is a post-fixed point.

Every chain in $\mathcal{F}$ has an upper bound in $\mathcal{F}$; in fact it has a least upper bound. Here is the argument. Let C be a chain in $\mathcal{F}$. Then $\sum C$ is consistent by Proposition 7.2 part 2. Suppose $w \in C$; then $w \leq_k \sum C$, so by monotonicity, $\varphi^e_{\mathcal{B}}(w) \leq_k \varphi^e_{\mathcal{B}}(\sum C)$. Since $w \in C \subseteq \mathcal{F}$, $w \leq_k \varphi^e_{\mathcal{B}}(w)$, so it follows that $w \leq_k \varphi^e_{\mathcal{B}}(\sum C)$ for every $w \in C$, and consequently $\sum C \leq_k \varphi^e_{\mathcal{B}}(\sum C)$. Then $\sum C \in \mathcal{F}$, so C has an upper bound in $\mathcal{F}$.

The ordering $\leq_k$, restricted to $\mathcal{F}$, is still a partial order, so $\mathcal{F}$ is a partially ordered set in which each chain has an upper bound. By Zorn's Lemma each member of $\mathcal{F}$ extends to a maximal member of $\mathcal{F}$. Finally we show that all consistent fixed points of $\varphi^e_{\mathcal{B}}$ are in $\mathcal{F}$, and that maximal members of $\mathcal{F}$ are themselves consistent fixed points. Then every maximal member of $\mathcal{F}$ will be a consistent maximal fixed point.

If a is a consistent fixed point of $\varphi^e_{\mathcal{B}}$ then, trivially, $a \leq_k \varphi^e_{\mathcal{B}}(a)$, so $a \in \mathcal{F}$. Suppose m is any maximal member of $\mathcal{F}$. By definition it is consistent. Also $m \leq_k \varphi^e_{\mathcal{B}}(m)$ so by monotonicity, $\varphi^e_{\mathcal{B}}(m) \leq_k \varphi^e_{\mathcal{B}}(\varphi^e_{\mathcal{B}}(m))$, and by the preservation of consistency property, $\varphi^e_{\mathcal{B}}(m)$ is consistent. Then $\varphi^e_{\mathcal{B}}(m) \in \mathcal{F}$. Since m is maximal in $\mathcal{F}$ and $m \leq_k \varphi^e_{\mathcal{B}}(m)$, it follows that $m = \varphi^e_{\mathcal{B}}(m)$, so m is a fixed point. □

9.2 Intrinsic Fixed Points

Kripke singled out an interesting subclass of fixed points for particular attention, the *intrinsic* ones, which are those compatible with every fixed point. An equivalent characterization also appeared independently in Manna and Shamir (1977). Here is a bilattice generalization that agrees with these on the consistent part of $\mathcal{FOUR}$.

Definition 9.4 (Intrinsic Fixed Point) *A fixed point v of $\varphi^e_{\mathcal{B}}$ is* intrinsic *if $v \oplus w$ is consistent, for every consistent fixed point w of $\varphi^e_{\mathcal{B}}$.*

There are intrinsic fixed points, and we have already seen one.

Proposition 9.5 *The smallest fixed point of $\varphi^e_{\mathcal{B}}$ is intrinsic, and every intrinsic fixed point is consistent.*

Proof Let v be the smallest fixed point with respect to $\leq_k$. If w is any consistent fixed point then $v \leq_k w$ so $v \oplus w = w$ which is consistent. Hence v is intrinsic.

Let w be any intrinsic fixed point. Since the smallest fixed point v is consistent, $v \oplus w$ must be consistent because w is intrinsic. But $v \oplus w$ is w, hence w is consistent.□

There is a direct connection between intrinsic fixed points and the consistent maximal ones of Section 9.1.

Proposition 9.6 *A fixed point v of $\varphi^e_{\mathcal{B}}$ is intrinsic if and only if $v \leq_k m$ for every consistent maximal fixed point if and only if $v \leq_k \prod \mathcal{M}$ where $\mathcal{M}$ is the set of all consistent maximal fixed points.*

Proof The if direction and the only if direction have different arguments.

1. Let v be an intrinsic fixed point, and m be a consistent maximal fixed point; we show $v \leq_k m$. Since v is intrinsic, $v \oplus m$ is consistent. Since $v \leq_k v \oplus m$, by monotonicity $\varphi^e_{\mathcal{B}}(v) \leq_k \varphi^e_{\mathcal{B}}(v \oplus m)$. Then $v \leq_k \varphi^e_{\mathcal{B}}(v \oplus m)$ since v is a fixed point. Similarly $m \leq_k \varphi^e_{\mathcal{B}}(v \oplus m)$, so $v \oplus m \leq_k \varphi^e_{\mathcal{B}}(v \oplus m)$. By Proposition 9.3 there is a consistent maximal fixed point m' above $v \oplus m$, that is, $v \oplus m \leq_k m'$. Then $m \leq_k v \oplus m \leq_k m'$. Since m is maximal, $m = m'$, and so $v \leq_k v \oplus m \leq_k m$.
2. Let v be a fixed point that is below every consistent maximal fixed point; we show v is intrinsic. Let w be an arbitrary consistent fixed point of $\varphi^e_{\mathcal{B}}$; we show $v \oplus w$ is consistent. Since w is a fixed point, $w \leq_k \varphi^e_{\mathcal{B}}(w)$, so by Proposition 9.3 $w \leq_k m$ for some consistent maximal fixed point m. $v \leq_k m$ by hypothesis. Then $v \oplus w \leq_k m \oplus m = m$, so also $-m \leq_k -(v \oplus w)$. Since m is consistent, $m \leq_k -m$. Then $(v \oplus w) \leq_k m \leq_k -m \leq_k -(v \oplus w)$, so $v \oplus w$ is consistent. □

In bilattices there is a largest as well as a smallest fixed point for truth revision operators. Confined to the consistent part there is no largest, though there are maximal ones. Confined further to the intrinsic fixed points there is a largest again, as was noted in Kripke (1975). The following includes some ideas from Fitting (1986) as well.

Proposition 9.7 *There is a largest intrinsic fixed point of $\varphi^e_{\mathcal{B}}$. It is $\sum \mathcal{I}$ where $\mathcal{I}$ is the set of all intrinsic fixed points. An ordinal sequence constructed using $\varphi^e_{\mathcal{B}}$ as sketched in the proof of Proposition 5.2, starting at $\prod \mathcal{M}$, where $\mathcal{M}$ is the set of consistent maximal fixed points, proceeds downward in the $\leq_k$ ordering and converges to the largest intrinsic fixed point.*

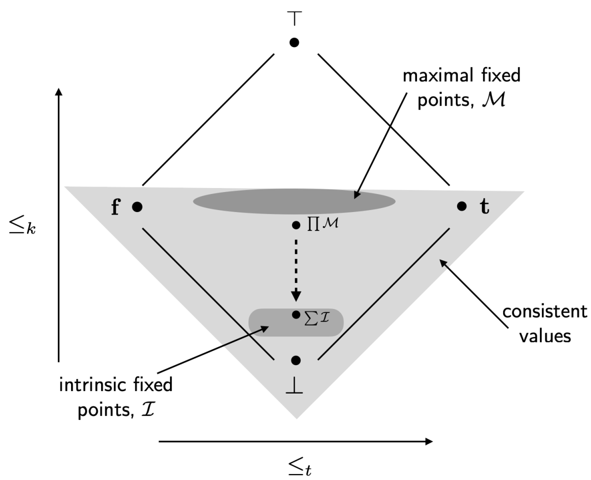

Fig. 5: Maximal and Intrinsic Fixed Points

Proof Figure 5 shows the overall relationships between the various kinds of fixed points of $\varphi^e_{\mathcal{B}}$ that will be established here. Keep in mind that it is highly schematic. The bilattice with conflation that is shown is a space of valuations; only the extremal values are shown explicitly. All our work is confined to the consistent part of it, which is shown lightly shaded. Two subsets of the consistent valuations are shown shaded darker. $\mathcal{I}$ is the set of intrinsic fixed points, which includes the least fixed point. $\mathcal{M}$ is the set of consistent maximal fixed points.

First we show that $\sum \mathcal{I} \le_k \prod \mathcal{M}$. Let v be an arbitrary member of $\mathcal{I}$. Since v is an intrinsic fixed point of $\varphi^e_{\mathcal{B}}$ then $v \le_k \prod \mathcal{M}$ by Proposition 9.6. Since v was arbitrary, $\sum \mathcal{I} \le_k \prod \mathcal{M}$.

Next we show that if v is a fixed point of $\varphi^e_{\mathcal{B}}$ and $\sum \mathcal{I} \le_k v \le_k \prod \mathcal{M}$ then v is in fact the largest intrinsic fixed point and $v = \sum \mathcal{I}$. Very simply, v is intrinsic by Proposition 9.6, and v is largest because, if w is any intrinsic fixed point, $w \in \mathcal{I}$ and so $w \le_k \sum \mathcal{I} \le_k v$. Finally since v is intrinsic, $v \in \mathcal{I}$ so $v \le_k \sum \mathcal{I}$, and so $v = \sum \mathcal{I}$.

Finally we show a fixed point between $\sum \mathcal{I}$ and $\prod \mathcal{M}$ exists. This can be done in two different ways, and we sketch both.

For the first way, $\sum \mathcal{I}$ is a post-fixed point of $\varphi^e_{\mathcal{B}}$ by the following argument. For every $w \in \mathcal{I}$ we have $w \le_k \sum \mathcal{I}$ so by monotonicity, $\varphi^e_{\mathcal{B}}(w) \le_k \varphi^e_{\mathcal{B}}(\sum \mathcal{I})$, and so $w \le_k \varphi^e_{\mathcal{B}}(\sum \mathcal{I})$ since w is a fixed point. Since w was arbitrary, $\sum \mathcal{I} \le_k \varphi^e_{\mathcal{B}}(\sum \mathcal{I})$. Then by Proposition 5.2 part 2, there is a fixed point that is least above $\sum \mathcal{I}$ and, using Corollary 5.4, it is easy to show it is below $\prod \mathcal{M}$.

For the second way, $\prod \mathcal{M}$ is a pre-fixed point of $\varphi^e_{\mathcal{B}}$, with a proof dual to that for $\sum \mathcal{I}$ above. Using Proposition 5.2 part 3, there is a fixed point that is greatest below $\prod \mathcal{M}$ and, using Corollary 5.4 dualized, one shows it is above $\sum \mathcal{I}$.

The first of the two proofs just sketched involves an ordinal sequence that, ostensibly, moves upward from $\sum \mathcal{I}$, but in fact it is actually constant. The second argument uses an ordinal sequence that proceeds downward from $\prod \mathcal{M}$, approximating to $\sum \mathcal{I}$ from above. This is suggested by the dashed arrow in the figure. □

9.3 Alternating Fixed Points

A very different class of fixed points was introduced in Yablo (1985), making use of *alternating* fixed points, Definition 5.5. These are closely related to the stable model semantics of logic programming, Gelfond and Lifschitz (1988), Fine (1989), and to well-founded semantics, Van Gelder (1989), Van Gelder, Ross, and Schlipf (1988, 1991). The treatment here is derived from Fitting (1993) on the logic programming side, and Fitting (1997) on the philosophical logic side.

A truth revision operator is monotone with respect to the $\leq_k$ ordering, but not in the $\leq_t$ ordering because of negation. It is possible, however, to localize this problem. Sentences of language $\mathcal{L}$ can have their negations pushed in to the atomic level using the various De Morgan laws. If this has been done we can think of occurrences of $\neg\mathsf{T}(x)$ as if they were occurrences of a new atom, a falsehood atom, that can behave independently of $\mathsf{T}(x)$.

Terminology Convention 9.8 For this subsection only, a sentence must be in negation normal form, and so all occurrences of the negation symbol are at the atomic level.

The following introduces a mapping taking two valuations as input and producing what we call a *pseudo valuation* as output, where this independently maps positive and negative literals involving T to bilattice values.

Definition 9.9 (Pseudo-Valuations) $\mathcal{A}(\neg\mathsf{T})$ *is the set of negated atomic sentences in the language* $\mathcal{L}$ *of the form* $\neg\mathsf{T}(0^{++\cdots+})$. *The map* $\triangle : \mathcal{B}^{\mathcal{A}(\mathsf{T})}\times\mathcal{B}^{\mathcal{A}(\mathsf{T})} \to \mathcal{B}^{\mathcal{A}(\mathsf{T})\cup\mathcal{A}(\neg\mathsf{T})}$ *is defined by the following conditions (where we write* $\triangle$ *in infix position):*

$$(v_1 \triangle v_2)(\mathsf{T}(0^{++\cdots+})) = v_1(\mathsf{T}(0^{++\cdots+}))$$
$$(v_1 \triangle v_2)(\neg\mathsf{T}(0^{++\cdots+})) = \neg v_2(\mathsf{T}(0^{++\cdots+})).$$

$(v_1 \triangle v_2)$ *is the* pseudo-valuation *of valuations* v_1 *and* v_2.

We had various ways of extending valuations to all formulas, and it is the same for pseudo-valuations. The following is the counterpart of Definitions 6.2 and 6.4.

Definition 9.10 (Pseudo-Valuation Extension) *A* pseudo-valuation extension *is a mapping* $\triangle^e : \mathcal{B}^{\mathcal{A}(\mathsf{T})}\times\mathcal{B}^{\mathcal{A}(\mathsf{T})} \to \mathcal{B}^{\mathcal{L}}$ *(where* $\mathcal{L}$ *is here the set of sentences in negation normal form). The* monotonicity/anti-monotonicity property *holds for* $\triangle^e$ *if:*

1. *k Monotonicity in Both Inputs:* $v_1 \leq_k v_2$ *and* $w_1 \leq_k w_2$ *implies* $(v_1 \triangle^e w_1) \leq_k (v_2 \triangle^e w_2)$.

2. *t Monotonicity in First Input:* $v_1 \le_t v_2$ *implies* $(v_1 \triangle^e w) \le_t (v_2 \triangle^e w)$.
3. *t Anti-Monotonicity in Second Input:* $w_1 \le_t w_2$ *implies* $(v \triangle^e w_1) \ge_t (v \triangle^e w_2)$.

Both strong Kleene and weak Kleene provide obvious pseudo-valuation extensions. For instance, Definition 8.7 for weak Kleene valuations, becomes the following, in bilattice $\mathcal{B}$.

1. For L being an atomic sentence or its negation, involving T, $(v_1 \triangle^w v_2)(L) = (v_1 \triangle v_2)(L)$.
2. For numerals a, b, and c:
   a. $(v_1 \triangle^w v_2)(A(a,b,c)) = v_1^w(A(a,b,c)) = v_2^w(A(a,b,c))$;
   b. $(v_1 \triangle^w v_2)(M(a,b,c)) = v_1^w(M(a,b,c)) = v_2^w(M(a,b,c))$;
   c. $(v_1 \triangle^w v_2)(a = b) = v_1^w(a = b) = v_2^w(a = b)$.
3. $(v_1 \triangle^w v_2)(X \wedge Y) = (v_1 \triangle^w v_2)(X) \wedge^w (v_1 \triangle^w v_2)(Y)$
4. $(v_1 \triangle^w v_2)(X \vee Y) = (v_1 \triangle^w v_2)(X) \vee^w (v_1 \triangle^w v_2)(Y)$
5. There is no separate negation case since formulas are in negation normal form.
6. $(v_1 \triangle^w v_2)(\forall x F(x)) = \bigwedge^w \{(v_1 \triangle^w v_2)(F(t)) \mid t \text{ is a numeral}\}$
7. $(v_1 \triangle^w v_2)(\exists x F(x)) = \bigvee^w \{(v_1 \triangle^w v_2)(F(t)) \mid t \text{ is a numeral}\}$

If language $\mathcal{L}$ is modal, we add the following, where $\mathcal{F} = \langle \mathcal{G}, \mathcal{R} \rangle$ is a modal frame.

8. $(v_1 \triangle^w v_2)(\Box X)(\Gamma) = \bigwedge^w \{(v_1 \triangle^w v_2)(X)(\Delta) \mid \Delta \in \mathcal{G} \text{ and } \Gamma \mathcal{R} \Delta\}$
9. $(v_1 \triangle^w v_2)(\Diamond X)(\Gamma) = \bigvee^w \{(v_1 \triangle^w v_2)(X)(\Delta) \mid \Delta \in \mathcal{G} \text{ and } \Gamma \mathcal{R} \Delta\}$

Similarly for strong Kleene. We omit details. Pseudo-valuation extensions of strong Kleene and weak Kleene both meet the monotonicity/anti-monotonicity conditions. This is straighhtforward to check, and we leave it to the reader. The status of pseudo-valuation extensions for supervaluations has not been investigated.

The following is the counterpart of Definition 6.3, for Truth Revision Operators.

Definition 9.11 (Two Input Truth Revision Operator) *Let* $\triangle$ *be a pseudo-valuation and* $\triangle^e$ *be an extension. Then* $\psi^e_{\mathcal{B}}(v_1, v_2) : \mathcal{B}^{\mathcal{A}(\mathsf{T})} \times \mathcal{B}^{\mathcal{A}(\mathsf{T})} \to \mathcal{B}^{\mathcal{A}(\mathsf{T})}$ *is the mapping such that for* T *valuations* v_1, v_2 *in* $\mathcal{B}$,

$$\psi^e_{\mathcal{B}}(v_1, v_2)(\mathsf{T}(\ulcorner X \urcorner)) = (v_1 \triangle^e v_2)(X) \text{ for every negation normal form sentence } X.$$

The following is an easily proved counterpart to Proposition 6.5.

Proposition 9.12 *If the* $\triangle^e$ *monotonicity/anti-monotonicity property holds (Definition 9.10). Then the two input truth revision operator* $\psi^e_{\mathcal{B}}$ *has the following properties.*

1. *k Monotonicity in Both Inputs: if* $v_1 \le_k v_2$ *and* $w_1 \le_k w_2$ *then* $\psi^e_{\mathcal{B}}(v_1, w_1) \le_k \psi^e_{\mathcal{B}}(v_2, w_2)$.
2. *t Monotonicity in First Input: if* $v_1 \le_t v_2$ *then* $\psi^e_{\mathcal{B}}(v_1, w) \le_t \psi^e_{\mathcal{B}}(v_2, w)$.
3. *t Anti-Monotonicity in Second Input: if* $w_1 \le_t w_2$ *then* $\psi^e_{\mathcal{B}}(v, w_1) \ge_t \psi^e_{\mathcal{B}}(v, w_2)$.

The original motivation comes from logic programming. Make a guess at negative values, somehow choosing a valuation v to tell us how occurrences of $\neg\mathsf{T}(t)$ behave. Relative to this guess, the least fixed point of $\lambda x.\psi^e_{\mathcal{B}}(x, v)$ tells a plausible way that positive $\mathsf{T}(t)$ atoms could then behave. If this turns out to be the guess we began with, then v was a good guess. Here the $\leq_t$ ordering plays a fundamental role. This suggest the following analog of Truth Revision Operators, Definition 6.3.

Definition 9.13 (Derived Operator) *The* derived operator *of the two input truth revision operator $\psi^e_{\mathcal{B}}$ is the single input function $\widehat{\psi^e_{\mathcal{B}}}$ where $\widehat{\psi^e_{\mathcal{B}}}(v)$ is the smallest fixed point, in the $\leq_t$ ordering, of the function $\lambda x.\psi^e_{\mathcal{B}}(x, v)$.*

Proposition 9.14 *Assuming monotonicity/anti-monotonicity for $\triangle^e$, the function $\widehat{\psi^e_{\mathcal{B}}}$ is well-defined, is monotonic in $\leq_k$, and anti-monotonic in $\leq_t$.*

Proof Assume $\psi^e_{\mathcal{B}}$ has the properties specified in Proposition 9.12.

$\widehat{\psi^e_{\mathcal{B}}}$ is well-defined. The smallest fixed point required in the definition of $\widehat{\psi^e_{\mathcal{B}}}$ exists because $\psi^e_{\mathcal{B}}$ is monotonic in its first input under $\leq_t$.

Monotonicity in the $\leq_k$ ordering. Assume $v_1 \leq_k v_2$. We define two ordinal sequences F and G by transfinite recursion. Both sequences consist of T valuations, members of $\mathcal{B}^{\mathcal{A}(\mathsf{T})}$. Also $\mathbf{f}$ is the T valuation that takes the value $\mathbf{f} \in \mathcal{B}$ on every input, $\alpha + 1$ is an arbitrary successor ordinal, and λ is an arbitrary limit ordinal.

$$
\begin{aligned}
F(0) &= \mathbf{f} & G(0) &= \mathbf{f} \\
F(\alpha + 1) &= \psi^e_{\mathcal{B}}(F(\alpha), v_1) & G(\alpha + 1) &= \psi^e_{\mathcal{B}}(F(\alpha), v_2) \\
F(\lambda) &= \bigvee\{F(\alpha) \mid \alpha < \lambda\} & G(\lambda) &= \bigvee\{G(\alpha) \mid \alpha < \lambda\}
\end{aligned}
$$

From the proof of Proposition 5.2, both ordinal sequences are increasing, and both are constant from some point on, the first settling on the least fixed point of $\lambda x.\psi^e_{\mathcal{B}}(x, v_1)$ in the $\leq_t$ ordering, and hence to $\widehat{\psi^e_{\mathcal{B}}}(v_1)$, and the second to $\widehat{\psi^e_{\mathcal{B}}}(v_2)$. To prove $\widehat{\psi^e_{\mathcal{B}}}(v_1) \leq_k \widehat{\psi^e_{\mathcal{B}}}(v_2)$ we show that, for each ordinal α, $F(\alpha) \leq_k G(\alpha)$. $F(\alpha) \leq G(\alpha)$ is trivially true for $\alpha = 0$. If it is true for all $\alpha < \lambda$, a limit ordinal, it is true for λ using infinitary interlacing. This leaves the successor step. Suppose $F(\alpha) \leq_k G(\alpha)$. Then $F(\alpha + 1) = \psi^e_{\mathcal{B}}(F(\alpha), v_1) \leq_k \psi^e_{\mathcal{B}}(G(\alpha), v_2) = G(\alpha + 1)$ using k monotonicity in both inputs.

Anti-monotonicity in the $\leq_t$ ordering. Assume $v_1 \leq_t v_2$; show $\widehat{\psi^e_{\mathcal{B}}}(v_2) \leq_t \widehat{\psi^e_{\mathcal{B}}}(v_1)$. $\psi^e_{\mathcal{B}}$ is anti-monotonic in its second argument under $\leq_t$, so $\psi^e_{\mathcal{B}}(\widehat{\psi^e_{\mathcal{B}}}(v_1), v_2) \leq_t \psi^e_{\mathcal{B}}(\widehat{\psi^e_{\mathcal{B}}}(v_1), v_1)$. Then $\psi^e_{\mathcal{B}}(\widehat{\psi^e_{\mathcal{B}}}(v_1), v_2) \leq_t \widehat{\psi^e_{\mathcal{B}}}(v_1)$ because $\widehat{\psi^e_{\mathcal{B}}}(v_1)$ is a fixed point of $\lambda x.\psi^e_{\mathcal{B}}(x, v_1)$. But then $\widehat{\psi^e_{\mathcal{B}}}(v_1)$ is a pre-fixed point of $\lambda x.\psi^e_{\mathcal{B}}(x, v_2)$ in the $\leq_t$ ordering, so there is a fixed point of $\lambda x.\psi^e_{\mathcal{B}}(x, v_2)$, call it a, below $\widehat{\psi^e_{\mathcal{B}}}(v_1)$. Since $\widehat{\psi^e_{\mathcal{B}}}(v_2)$ is the *least* fixed point of $\lambda x.\psi^e_{\mathcal{B}}(x, v_2)$, then $\widehat{\psi^e_{\mathcal{B}}}(v_2) \leq_t a \leq_t \widehat{\psi^e_{\mathcal{B}}}(v_1)$. □

Definition 9.15 (Stable Fixed Points) *Let $\psi^e_{\mathcal{B}}$ be a two input truth revision operator, and let $\widehat{\psi^e_{\mathcal{B}}}$ be the corresponding derived operator. We call fixed points of $\widehat{\psi^e_{\mathcal{B}}}$* stable *e fixed points.*

Monotonicity of derived operators in the $\leq_k$ ordering tells us that stable fixed points exist. Assuming reasonable conditions, stable fixed points are a subclass of the class of fixed points we have been looking at throughout.

Definition 9.16 (Connected Condition) *We say pseudo-valuation extension e and valuation extension e are* connected *if, for each* T *valuation v and each sentence X (in negation normal form, remember),* $(v \triangle^e v)(X) = v^e(X)$.

Both strong Kleene and weak Kleene have connected valuation and pseudo-valuation extensions. This is easy, and we leave it to the reader.

Proposition 9.17 *Assume the monotonicity/anti-monotonicity conditions hold, and pseudo-valuation extension e and valuation extension e are connected. Then every fixed point of* $\widehat{\psi^e_{\mathcal{B}}}$ *is also a fixed point of* $\varphi^e_{\mathcal{B}}$ *(Definition 6.3). That is, every stable fixed point is a fixed point in the sense of Kripke.*

Proof It follows easily from the definition of being connected that, for each T valuation v, $\varphi^e_{\mathcal{B}}(v) = \psi^e_{\mathcal{B}}(v, v)$. Also for each valuation v, $\widehat{\psi^e_{\mathcal{B}}}(v)$ is a fixed point of $\lambda x.\psi^e_{\mathcal{B}}(x, v)$ so $\psi^e_{\mathcal{B}}(\widehat{\psi^e_{\mathcal{B}}}(v), v) = \widehat{\psi^e_{\mathcal{B}}}(v)$. Now suppose s is any fixed point of $\widehat{\psi^e_{\mathcal{B}}}$. Then $\varphi^e_{\mathcal{B}}(s) = \psi^e_{\mathcal{B}}(s, s) = \psi^e_{\mathcal{B}}(\widehat{\psi^e_{\mathcal{B}}}(s), s) = \widehat{\psi^e_{\mathcal{B}}}(s) = s$. □

Stable fixed points are thus among the family of Kripke style fixed points, but they are a *proper* subset. With strong Kleene, and using a Kripke style operator, fixed points can give the truth teller any value from the bilattice. This is not the case with stable fixed points. The definition uses the smallest fixed point in the truth ordering when evaluating the derived operator, so an explicit bias towards falsehood has been introduced, and in a stable strong Kleene fixed point the truth teller can only be **f**.

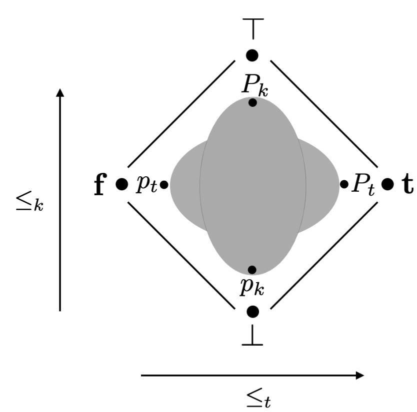

Fig. 6: Stable Fixed points

We know from Proposition 9.14 that the derived operator $\widehat{\psi^e_{\mathcal{B}}}$ is monotone under $\leq_k$ and anti-monotone under $\leq_t$. Because of monotonicity it has least and greatest fixed points under $\leq_k$, call them p_k and P_k. Then, of course, all stable fixed points are between p_k and P_k in the information ordering, with the extremes included. Further, because of anti-monotonicity under $\leq_t$ the operator $\widehat{\psi^e_{\mathcal{B}}}$ has an extremal alternating fixpoint pair by Proposition 5.6, call the members of the pair p_t and P_t. While these constitute an alternating fixpoint pair, they are not themselves fixed points. But each fixed point can be thought of as an alternating fixpoint pair, with both members of the pair being the same, and so all stable fixed points are strictly between p_t and P_t in the truth ordering. The space of stable fixed points is shown schematically in Figure 6. Further, the four extremal values are tightly connected; by Proposition 5.9: $P_t \otimes p_t = p_k$, $P_t \oplus p_t = P_k$, $P_k \wedge p_k = p_t$, $P_k \vee p_k = P_t$.

In Fitting (1997) the notion of intrinsic (Section 9.2) was relativized to stable fixed points. A stable fixed point v is *stable intrinsic* if $v \oplus w$ is consistent for every consistent stable fixed point valuation w. There are stable intrinsic fixed points, but essentially no investigation of their properties has been made.

10 Conclusion

Several versions of theories of truth have been examined, in very general settings. Still, there could be more. For example, Kripke (1975, p. 706) made some interesting suggestions.

> The construction could be generalized so as to allow more notation in $\mathcal{L}$ than just first-order logic. For example, we could have a quantifier meaning "for uncountably many x," a "most" quantifier, a language with infinite conjunctions, etc. There is a fairly canonical way, in the Kleene style, to extend the semantics of such quantifiers and connectives so as to allow truth-value gaps, but we will not give details.

As far as I know, this has not been followed up in the original three valued setting. It certainly has not been using bilattice generalizations. It seems to be an open invitation still today. Kripke also mentioned including a modal operator. In the present treatment this is simply working with a bilattice of a special form, as in Example 4.9. One might carry this further and allow multiple knowledge modalities. Perhaps even communication between agents could be modeled. Then, of course, we could model the quirks of self referential truth that arise in the discussions of a group of philosophers. (This is not an entirely frivolous suggestion.) We hope the present paper enables possibilities, and does not represent the end of the discussion.

It has been suggested that some detailed comments on the philosophical implications of the work presented here might be given, but I have chosen not to do this. I feel I have presented a rich set of tools. It is hoped that philosophical investigators will find some of the tools useful. But I am a theorist, and applications are not really my job.

References

Arieli, O. & Avron, A. (1998). The value of four values. *Artificial Intelligence*, *102*, 97–141.

Avron, A. (1996). The structure of interlaced bilattices. *Mathematical Structures in Computer Science*, *6*(3), 287–299.

Belnap, N. D., Jr. (1977). A useful four-valued logic. In J. M. Dunn & G. Epstein (Eds.), *Modern uses of multiple-valued logic* (pp. 5–37). Springer Netherlands.

Bochvar, D. A. (1938). Ob odnom trechznacnom iscislenii i ego primenenii k analizu paradoksov klassiceskogo rassirennogo funkcional'nogo iscislenija. *Matematiceskij Sbornik*, *4*, 287–308. English translation in Bochvar, 1981.

Bochvar, D. A. (1981). On a three-valued logical calculus and its application to the analyses of the paradoxes of the classical extended functional calculus. *History and Philosophy of Logic*, *2*(87-112). Merrie Bergmann translator.

Cobreros, P. (2013). Vagueness: Subvaluationism. *Philosophy Compass*, *8*(5), 472–485.

Davey, B. (2013). The product representation theorem for interlaced pre-bilattices: Some historical remarks. *Algebra Universalis*. (In press.)

Ferguson, T. (2015). Cut-down operations on bilattices. In *2015 ieee international symposium on multiple-valued logic* (pp. 24–29).

Fine, K. (1989). The justification of negation as failure. In J. E. Fenstad, I. T. Frolov, & R. Hilpinen (Eds.), *Logic, methodology and philosophy of science VIII* (pp. 263–301). Amsterdam: North-Holland.

Fitting, M. C. (1986). Notes on the mathematical aspects of Kripke's theory of truth. *Notre Dame Journal of Formal Logic*, *27*, 75–88.

Fitting, M. C. (1989). Bilattices and the theory of truth. *Journal of Philosophical Logic*, *18*, 225–256.

Fitting, M. C. (1990). Bilattices in logic programming. In G. Epstein (Ed.), *The twentieth international symposium on multiple-valued logic* (pp. 238–246). IEEE.

Fitting, M. C. (1991). Bilattices and the semantics of logic programming. *Journal of Logic Programming*, *11*, 91–116.

Fitting, M. C. (1993). The family of stable models. *Journal of Logic Programming*, *17*, 197–225.

Fitting, M. C. (1994). Kleene's three-valued logics and their children. *Fundamenta Informaticae*, *20*, 113–131.

Fitting, M. C. (1997). A theory of truth that prefers falsehood. *Journal of Philosophical Logic*, *26*, 477–500.

Fitting, M. C. (2006). Bilattices are nice things. In T. Bolander, V. Hendricks, & S. A. Pederrsen (Eds.), *Self-reference* (Chap. 3, pp. 53–77). Center for the Study of Language and Information.

Gargov, G. (1999). Knowledge, uncertaity and ignorance in logic: Bilattices and beyond. *Journal of Applied Non-Classical Logics*, *9*(2-3), 195–283.

Gelfond, M. & Lifschitz, V. (1988). The stable model semantics for logic programming. In R. Kowalski & K. Bowen (Eds.), *Proc. of the fifth logic programming symposium* (pp. 1070–1080). Cambridge, MA: MIT Press.

Ginsberg, M. L. (1988). Multivalued logics: A uniform approach to reasoning in artificial intelligence. *Computational Intelligence*, *4*, 265–316.

Kleene, S. C. (1950). *Introduction to metamathematics*. Princeton, NJ: D. Van Nostrand.

Kripke, S. (1975). Outline of a theory of truth. *The Journal of Philosophy*, *72*(19), 690–716. Reprinted in *New Essays on Truth and the Liar Paradox*, R. L. Martin, ed., Oxford (1983).

Manna, Z. & Shamir, A. (1977). The optimal approach to recursive programs. *Comm. ACM*, *20*, 824–831.

Martin, R. L. & Woodruff, P. W. (1975). On representing "true-in-L" in L. *Philosophia*, *5*, 217–221. Reprinted in *Recent Essays on Truth and the Liar paradox*, R. L. Martin editor, Oxford, 1984.

Pynko, A. P. (2000). Regular bilattices. *Journal of Applied Non-Classical Logics*, *10*, 93–111.

Smullyan, R. M. (1968). *First-order logic*. (Revised Edition, Dover Press, New York, 1994). Berlin: Springer-Verlag.

Szmuc, D. J. (2018). Track-down operations on bilattices. In R. Wille & M. Lukac (Eds.), *Proceedings of the 48th IEEE international symposium on multiple-valued logic* (pp. 74–79).

Teijeiro, P. (2020 forthcoming). Subvaluationism and classical recapture. *Logic Journal of the IGPL.*

Van Fraassen, B. (1966). Singular terms, truth-value gaps, and free logic. *Journal of Philosophy*, *63*, 481–485.

Van Gelder, A. (1989). The alternating fixpoint of logic programs with negation. In *Proc. 8th acm symp. on principles of database systems* (pp. 1–10). Philadelphia: ACM.

Van Gelder, A., Ross, K. A., & Schlipf, J. S. (1988). Unfounded sets and well-founded semantics for general logic programs. In *Proc. seventh symp. on principles of database systems* (pp. 221–230).

Van Gelder, A., Ross, K. A., & Schlipf, J. S. (1991). The well-founded semantics for general logic programs. *JACM*, *38*, 620–650.

Visser, A. (1984). Four valued semantics and the liar. *Journal of Philosophical Logic*, *13*, 181–212.

Yablo, S. (1985). Truth and reflection. *Journal of Philosophical Logic*, *14*, 297–349.