

The Family of Stable Models

Melvin Fitting

mlflc@cunyvm.cuny.edu

Dept. Mathematics and Computer Science

Lehman College (CUNY), Bronx, NY 10468

Depts. Computer Science, Philosophy, Mathematics

Graduate Center (CUNY), 33 West 42nd Street, NYC, NY 10036 *

January 5, 1993

Abstract

The family of all stable models for a logic program has a surprisingly simple overall structure, once two naturally occurring orderings are made explicit. In a so-called knowledge ordering based on degree of definedness, every logic program \mathcal{P} has a smallest stable model, $s_{\mathcal{P}}^k$ — it is the well-founded model. There is also a dual largest stable model, $S_{\mathcal{P}}^k$, which has not been considered before. There is another ordering based on degree of truth. Taking the meet and the join, in the truth ordering, of the two extreme stable models $s_{\mathcal{P}}^k$ and $S_{\mathcal{P}}^k$ just mentioned, yields the alternating fixed points of [29], denoted $s_{\mathcal{P}}^t$ and $S_{\mathcal{P}}^t$ here. From $s_{\mathcal{P}}^t$ and $S_{\mathcal{P}}^t$ in turn, $s_{\mathcal{P}}^k$ and $S_{\mathcal{P}}^k$ can be produced again, using the meet and join of the knowledge ordering. All stable models are bounded by these four valuations. Further, the methods of proof apply not just to logic programs considered classically, but to logic programs over any bilattice meeting certain conditions, and thus apply in a vast range of settings. The methods of proof are largely algebraic.

1 Introduction

Stable model semantics [17, 6] and well-founded model semantics [30, 31] have provided valuable insights on the meaning of logic programs. While much is known about these semantics, it has not been made clear just how elegant a structure the family of stable models has. The delineation of this structure is the primary aim of this paper. We do not attempt to state our results rigorously at this point, but wait until Section 9, after appropriate terminology and concepts have been introduced. We can give the general flavor, however.

As originally formulated, stable model semantics was classical, two-valued. Under these circumstances not every program has a stable model. Three-valued, or partial model semantics has had an extensive development for logic programs generally, [7, 20, 21]. Soon Przymusiński extended the notion of stable model to allow three-valued, or partial, stable models, [22, 23], and showed every program has at least one partial stable model, and the well-founded model is the smallest among them, in a natural ordering, [24]. Once one has made the transition from classical to partial models allowing incomplete information, it is a small step to also allow models admitting inconsistent information. Doing so provides a natural framework for the semantic understanding of logic

*Research partly supported by NSF Grants CCR-8901489 and CCR-9104015.

programs that are distributed over several sites, with possibly conflicting information coming from different places, [10]. We make the extension to such a setting in this paper, as part of a much more general investigation.

One can formally think of allowing incompleteness or inconsistency of information in logic programs in the following way. Take a valuation to be a mapping from ground atoms to *sets* of truth values. The values $\{true\}$ and $\{false\}$ correspond to the usual notions of truth and falsity, while \emptyset and $\{false, true\}$ correspond to no information and inconsistent information, respectively. These are the four truth values of a logic due to Belnap, [3], one that contains the Kleene strong three-valued logic as a sublogic. Now the space of valuations has two natural orderings. One is on the ‘degree of truth.’ In this ordering an increase means that one or more ground atoms loses an assignment of *false*, or gains an assignment of *true*. The other ordering is on the ‘degree of knowledge.’ In this ordering an increase means at least one ground atom gains an assignment of *true* or *false* that it did not have, without losing any value it did have. As we will see, these two orderings are intimately connected. Now, we will show that the family of stable models for a program is bounded from above and below, in the knowledge ordering, by a biggest and a smallest stable model. The smallest is the well-founded model; this is Przymusiński’s result. The existence of a largest stable model is new, and could not have been formulated without moving to the setting of Belnap’s logic. Also the family of stable models is bounded from above and below in the truth ordering by two valuations that are not necessarily stable models themselves, but which are extremal in a certain sense. This, essentially, is the ‘alternating fixpoint construction’ of [29]. Finally, we will show that the smallest and biggest stable models in the knowledge ordering, and the extremal valuations in the truth ordering, are closely connected, and can be ‘calculated’ from each other. This is new.

The results just sketched actually hold in a much broader setting than was indicated. Establishing them in this broader setting takes no more work than it does in the classical context. In addition, working in the extended context tends to strip away some of the unnecessary detail and make the underlying simplicity of the proofs stand out. We briefly outline the setting we have in mind.

The four-valued logic of Belnap is the simplest example of a non-trivial *bilattice*, a notion introduced by Matt Ginsberg in [19]. We have made use of bilattices in previous work on the semantics of logic programming [11, 12, 13], and once again we find them the proper tool. Motivation will be supplied in more detail in Section 4. For now, we merely remark that bilattices arise naturally in two ways, which we sketch briefly.

First, instead of the all-or-nothing approach of classical logic, we might introduce a more general notion of evidence, such as a confidence factor. Then, further, we might separate the roles of positive and negative evidence. This leads directly to bilattices, and is the approach presented in [11, 12]. From this point of view, classical logic programming is based on there being only two kinds of evidence: totally convincing, and totally unconvincing.

Alternatively we might think of the underlying logic as being more general than two-valued. For instance, if we have a program divided among two sites, then an answer to a ground query should contain information on which sites answer yes, and so there are four natural truth values: neither, one, the other, both. We might even allow dependencies between sites. This gives rise to an algebraic structure of truth values [25, 15]. If, further, we admit the possibility of incomplete or inconsistent information about the behavior of these sites, we once again get the structure of a bilattice [14]. A summary of this, and the previous approach, will be presented in Section 4.

The work in this paper extends results previously presented in [13]. At the time the earlier paper was written, the real simplicity of the structure of stable models had not become clear, hence

the need for this fuller presentation. We remark again that moving from the classical setting to that of bilattices adds no complications. Indeed, we believe it makes the underlying simplicity of the results stand out by removing unessential details.

2 Belnap's logic, *FOUR*

In [3], Belnap introduced a logic intended to deal in a useful way with inconsistent or incomplete information. It is the simplest example of a non-trivial bilattice and it illustrates many of the basic ideas concerning them. In the next section we present the general notion of bilattice, and in the following one we show there are lots of them, and they arise in natural ways. But in fact, Belnap's logic is a good enough representative of the whole family so that, if you wish, on a first reading of this paper you could skip from the end of this section directly to Section 5, just assuming that every bilattice discussed is sufficiently like the Belnap one.

We can think of Belnap's values as *sets* of ordinary truth values. Suppose we write, simply, *true* for $\{\text{true}\}$ and *false* for $\{\text{false}\}$. Also we write \perp for \emptyset and think of it as indicating a lack of information. And we write \top for $\{\text{false}, \text{true}\}$, and think of it as indicating inconsistency.

The truth values of Belnap's logic have two natural orderings. One is the subset relation. For instance, $\perp = \emptyset \subseteq \{\text{false}\} = \text{false}$. If $A \subseteq B$, in a natural sense B represents more information than A , so this is referred to as a *knowledge* ordering, and is denoted by \leq_k . Thus $\perp \leq_k \text{false}$. The other ordering is on the *degree of truth*, and is denoted \leq_t . Here $A \leq_t B$ if ' B is at least as true as A is, and A is at least as false as B is.' More precisely, $A \leq_t B$ if $A \cap \{\text{true}\} \subseteq B \cap \{\text{true}\}$ and $B \cap \{\text{false}\} \subseteq A \cap \{\text{false}\}$. Under this ordering, $\text{false} \leq_t \perp$. The two orderings are represented in the double Hasse diagram shown in Figure 1.

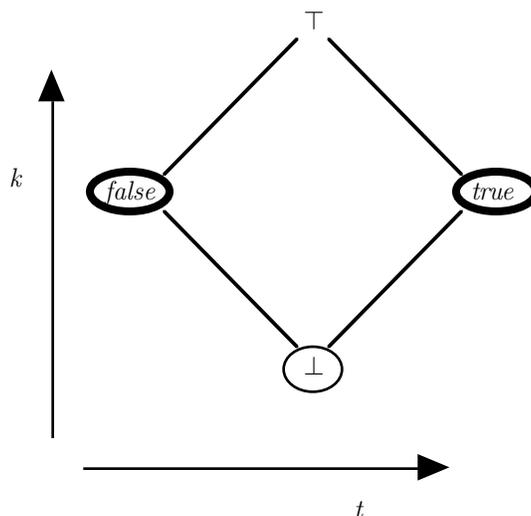


Figure 1: The logic *FOUR*

Both \leq_k and \leq_t give *FOUR* a lattice structure. Meet and join under \leq_t are denoted \wedge and \vee ; they are natural generalizations of the usual conjunction and disjunction notions. Meet and join under \leq_k are denoted \otimes and \oplus . $x \otimes y$ is the most information x and y can agree on — we call \otimes

a *consensus* operator. Likewise $x \oplus y$ simply combines the knowledge represented by x with that represented by y , without checking for consistency — we call \oplus a *gullability* operator. The four operations, \wedge , \vee , \otimes , and \oplus are intimately connected: all 12 distributive laws hold.

There is a natural notion of negation: flip the diagram from left to right, switching *false* and *true*, leaving \perp and \top alone. This is denoted \neg . There is also a natural symmetry vertically, switching \perp and \top , leaving *false* and *true* alone. This is denoted $-$ and is called *conflation*. Conflation and negation commute: $-\neg x = \neg - x$.

The role of conflation is less immediate than that of negation. Notice that the truth values left unchanged under conflation are *false* and *true*, the classical ones. These are marked with dark circles in Figure 1. The operations \neg , \wedge , and \vee , restricted to the values *false* and *true* reduce to those of classical logic. In addition, the values of *FOUR* that contain no more information than their conflation ($x \leq_k -x$) are *false*, *true*, and \perp , circled in Figure 1. It turns out that the operations \wedge , \vee , and \neg , restricted to *false*, *true*, and \perp are exactly those of Kleene’s strong three-valued logic. This way of distinguishing natural sublogics of *FOUR*, using conflation, is something that carries over to bilattices generally.

3 Bilattice definitions

Now we turn to the general notion of bilattice — keep *FOUR* in mind as a representative example. A bilattice is a lattice-like structure with two inter-related orderings. Loosely, one of the orderings represents degree of truth, the other represents degree of knowledge, as with *FOUR*. We sketch ways bilattices arise in the next section, after the formal definitions have been presented. One can make assumptions of various strengths concerning relationships between the orderings — in this way one obtains different kinds of bilattices. We will need rather strong assumptions for this paper.

Definition 3.1 A *pre-bilattice* is a structure $\langle \mathcal{B}, \leq_t, \leq_k \rangle$ where \mathcal{B} is a non-empty set and \leq_t and \leq_k are each partial orderings giving \mathcal{B} the structure of a lattice with a top and a bottom.

Remark As originally defined, in [19], both \leq_t and \leq_k were required to yield *complete* lattices. In various subsequent papers completeness has sometimes been assumed, sometimes not. Generally it has been assumed if an application involved quantifiers, but it was not assumed if algebraic properties of bilattices were being studied. We feel it is time to end this confusion. In this paper we do not assume completeness without saying so. Thus a pre-bilattice will specifically be called *complete* if its partial orderings give the structure of complete lattices. Similarly for interlaced and distributive bilattices below.

Definition 3.2 In a pre-bilattice $\langle \mathcal{B}, \leq_t, \leq_k \rangle$, meet and join under \leq_t are denoted \wedge and \vee , and meet and join under \leq_k are denoted \otimes and \oplus . Top and bottom under \leq_t are denoted *true* and *false*, and top and bottom under \leq_k are denoted \top and \perp . If the pre-bilattice is complete, infinitary meet and join under \leq_t are denoted \bigwedge and \bigvee , and infinitary meet and join under \leq_k are denoted \prod and \sum .

The operations \wedge and \vee should be thought of as generalizations of conjunction and disjunction in the classical space $\{\textit{false}, \textit{true}\}$. The operations \otimes , consensus, and \oplus , gullability, are not meaningful in the classical setting. One needs at least the structure of *FOUR* for them to arise. Now we define the various basic notions we will be using here.

Definition 3.3 A *distributive bilattice* is a pre-bilattice $\langle \mathcal{B}, \leq_t, \leq_k \rangle$ in which all 12 distributive laws connecting \wedge , \vee , \otimes and \oplus hold. An *infinitely distributive bilattice* is a complete pre-bilattice in which all infinitary, as well as all finitary, distributive laws hold.

An example of a distributive law is: $x \otimes (y \vee z) = (x \otimes y) \vee (x \otimes z)$. An example of an infinitary distributive law is: $x \otimes \bigwedge \{y_i \mid i \in S\} = \bigwedge \{x \otimes y_i \mid i \in S\}$. There is a weaker notion than distributivity that also has played a role in the semantics of logic programming [11, 12].

Definition 3.4 The pre-bilattice $\langle \mathcal{B}, \leq_t, \leq_k \rangle$ satisfies the *interlacing conditions* if each of the lattice operations, \wedge , \vee , \otimes , \oplus , is monotone with respect to both orderings. If the pre-bilattice is complete, it satisfies the *infinitary interlacing conditions* if each of the infinitary meet and join operations is monotone with respect to both orderings.

An example of an interlacing condition is: $x_1 \leq_t y_1$ and $x_2 \leq_t y_2$ implies $x_1 \otimes x_2 \leq_t y_1 \otimes y_2$. An example of an infinitary interlacing condition is: $x_i \leq_t y_i$ for each $i \in S$ implies $\prod \{x_i \mid i \in S\} \leq_t \prod \{y_i \mid i \in S\}$. A distributive bilattice meets the interlacing conditions as well. This is easily verified. But an infinitely distributive bilattice need not meet the infinitary interlacing conditions. We will explicitly assume such conditions, when needed.

As defined in [19], bilattices were required to have a negation operation, and the connection between the two orderings was a very simple one, expressed via negation. We find this too weak for our purposes, though negation still plays an important role here, along with the ‘dual’ operation of conflation. In the following, *bilattice* refers to any of the above: distributive bilattice, complete distributive bilattice, interlaced bilattice, or complete interlaced bilattice.

Definition 3.5 A bilattice has a *negation* if there is a mapping \neg that reverses the \leq_t ordering, leaves unchanged the \leq_k ordering, and $\neg\neg x = x$. Similarly a bilattice has a *conflation* if there is a mapping $-$ that reverses the \leq_k ordering, leaves unchanged the \leq_t ordering, and $--x = x$. If a bilattice has both, we say they *commute* if $--\neg x = \neg -x$ for all x .

Bilattices can have negations without conflations, conflations without negations, neither, or both. Examples are not hard to come by.

4 Bilattice constructions

Bilattices come up in natural ways — indeed the simplest non-trivial bilattice, *FOUR*, was already introduced in [3]. Two general, but different, constructions have appeared in the literature. We sketch them briefly here in order to give a feeling for bilattices and their possible applications.

The first construction comes from [19], and is completely general for distributive bilattices — that is, a representation theorem can be proved. Here we follow the presentation of [11].

Suppose we have two lattices, $\langle L_1, \leq_1 \rangle$ and $\langle L_2, \leq_2 \rangle$. Think of L_1 as the lattice of values we use when we measure the degree of belief, evidence, confidence, etc. that we have in a sentence. Likewise, think of L_2 as what we use when we measure the degree of doubt, counter-evidence, lack of confidence, etc. that we have against a sentence. For instance, L_1 and L_2 could both be just $\{false, true\}$, with $false \leq true$, where we have all-or-nothing judgments. Or each could be the unit interval, with the standard ordering. Or L_1 could be the unit interval and L_2 could be $\{false, true\}$, a choice that is appropriate for a theoretical science, where one has a degree of confidence in a sentence, but a single counter-experiment is enough to falsify. Or again, L_1 and L_2 could be sets of

experts, ordered by inclusion (the $\{false, true\}$ example can be thought of as arising from a single expert).

Now we define a structure $L_1 \odot L_2$ as follows. The structure is $\langle L_1 \times L_2, \leq_t, \leq_k \rangle$, where:

- $\langle x_1, x_2 \rangle \leq_t \langle y_1, y_2 \rangle$ if $x_1 \leq_1 y_1$ and $y_2 \leq_2 x_2$;
- $\langle x_1, x_2 \rangle \leq_k \langle y_1, y_2 \rangle$ if $x_1 \leq_1 y_1$ and $x_2 \leq_2 y_2$.

The idea is: knowledge goes up if evidence, both for and against, goes up; degree of truth goes up if evidence for goes up while evidence against goes down.

It is straightforward to verify that $L_1 \odot L_2$ will always be a bilattice meeting the interlacing conditions, and if L_1 and L_2 are complete, $L_1 \odot L_2$ will be a complete bilattice meeting the infinitary interlacing conditions as well. Further, if both L_1 and L_2 are distributive lattices, $L_1 \odot L_2$ will be a distributive bilattice, and if L_1 and L_2 are complete and infinitely distributive, $L_1 \odot L_2$ will be an infinitely distributive bilattice. This can be carried yet further. If $L_1 = L_2 = L$ then we are measuring belief and doubt the same way, and so a negation can be introduced into the bilattice $L \odot L$ in the obvious way: set $\neg \langle x, y \rangle = \langle y, x \rangle$. This is easily seen to meet the appropriate conditions for a negation operation. The intuition is straightforward: in passing from a member of $L \odot L$ to its negation we simply switch the roles of belief and doubt. Finally, suppose $L_1 = L_2 = L$, and further suppose that L has an order reversing involution (we denote the involute of x by \bar{x}). Such an operation is called a *de Morgan complement*. Then a conflation operation can be introduced into $L \odot L$ as well. Set $\neg \langle x, y \rangle = \langle \bar{y}, \bar{x} \rangle$. Once again it is easy to verify that this meets the conditions for a conflation operation, and conflation and negation will commute.

The idea behind this notion of conflation is a little more complicated than that behind negation. Think, for example, of L as being the unit interval, with members measuring degree of belief or doubt. Set $\bar{x} = 1 - x$, which is a de Morgan complement. Then $\neg \langle x, y \rangle = \langle 1 - y, 1 - x \rangle$, and so in passing from a member of $L \odot L$ to its conflation, we replace the degree of belief by the degree to which we did not doubt, and we replace the degree of doubt by the degree to which we did not believe. Many other examples of this nature can be given.

A *de Morgan lattice* is a distributive lattice with a de Morgan complement. Thus we have: if L is a de Morgan lattice, $L \odot L$ is a distributive bilattice with a negation and a conflation that commute.

The simplest non-trivial de Morgan lattice is $L = \{false, true\}$, of course, with $false \leq true$. The corresponding bilattice, $L \odot L$, is just *FOUR* all over again, with $\perp = \langle false, false \rangle$, $\top = \langle true, true \rangle$, $false = \langle false, true \rangle$, and $true = \langle true, false \rangle$. The bilattice *FOUR* has a long history. It was introduced in [4], outfitted with two orderings in [3], and recognized as a bilattice in [19].

The method just sketched for constructing bilattices is general for the distributive case. That is, one can show representation theorems along the following lines. If \mathcal{B} is a distributive bilattice, it is isomorphic to $L_1 \odot L_2$, where L_1 and L_2 are distributive lattices. If \mathcal{B} has a negation, L_1 and L_2 can be taken to be the same lattice. If \mathcal{B} has a negation and a conflation that commute, L_1 is a de Morgan lattice. Proofs of this can be found in [19, 11, 16]. In fact, the result can be looked at as a variation on the Polarities Theorem of Dunn, which goes back to his dissertation of 1966 (see [5]).

A bilattice with conflation and negation that commute with each other contains some natural sublogics within it. Suppose \mathcal{B} is such a bilattice. Call a member $x \in \mathcal{B}$ *consistent* if $x \leq_k \neg x$, and call x *exact* if $x = \neg x$. Just as with *FOUR*, the exact members of a bilattice always form a generalization of the classical truth values. In particular, they are closed under \wedge , \vee , and \neg . In the bilattice $L \odot L$, where L is the unit interval, for instance, the exact members are those of the form

$\langle x, y \rangle$ where $y = 1 - x$; thus they are the ones for which doubt exactly complements belief. The consistent members of a bilattice also are closed under \wedge , \vee , and \neg . In addition they are closed under \otimes , and under \oplus when applied to a set that is directed under \leq_k . In [9] these facts were used to generalize Kripke's theory of truth to bilattices, and in [12] they played a role in developing a semantics for logic programming. In the bilattice based on the unit interval, mentioned above, the consistent values are those $\langle x, y \rangle$ for which $x + y \leq 1$, and thus represent situations where belief and doubt are not contradictory.

There is a second way in which bilattices arise naturally. It was sketched in [19], and developed in detail in [14]. The idea is to think of the consistent members as 'approximations' to the exact values. This approach leads directly to Kleene's three-valued logic, and indirectly to *FOUR*, rather than the other way around, as above.

Suppose we have a lattice of truth values L . Think of these as the 'real' values we are interested in. In classical logic the intended lattice is $\{\text{false}, \text{true}\}$. If we have two independent experts, say A and B we have a more complicated lattice to work with, because the experts might disagree. In this case the natural lattice of truth values to work with is $\{\emptyset, \{A\}, \{B\}, \{A, B\}\}$, ordered by inclusion, where we think of a set of experts as a representation that those experts said "yes" to a query. Other lattices of truth values arise if we allow probabilities as truth values, and so on.

In general we may be uncertain about the status of a sentence; what truth value in L to assign to it. Our ignorance may not be total however, we may have some information to work with. In [14] approximations to a 'real' truth value were represented by intervals containing that truth value.

Definition 4.1 Let $a, b \in L$, where L is a lattice, and assume $a \leq_L b$. The *interval* $[a, b]$ is $\{x \in L \mid a \leq_L x \leq_L b\}$. $\mathcal{I}(L)$ is the set of all intervals in L .

Now we take the space of intervals and give it a bilattice-like structure.

Definition 4.2 Let L be a lattice. $\mathcal{K}(L)$ is the structure $\langle \mathcal{I}(L), \leq_t, \leq_k \rangle$ where, for $[a, b], [c, d] \in \mathcal{I}(L)$:

1. $[a, b] \leq_t [c, d]$ if $a \leq_L c$ and $b \leq_L d$;
2. $[a, b] \leq_k [c, d]$ if $[c, d] \subseteq [a, b]$.

The idea is, knowledge increases if the approximation interval narrows, and degree of truth increases if the interval shifts upward in the ordering of L .

This can be carried a little further. Suppose L also has an order-reversing involution, a de Morgan complement (again we denote the de Morgan complement of x by \bar{x}). Then a negation can be introduced in $\mathcal{K}(L)$ by:

$$\neg[a, b] = [\bar{b}, \bar{a}].$$

As an example, suppose we start with the lattice of classical truth values, $L = \{\text{false}, \text{true}\}$. Then $\mathcal{K}(L)$ has three members, two one-point intervals, $[\text{false}, \text{false}]$ and $[\text{true}, \text{true}]$, and the entire lattice, $[\text{false}, \text{true}]$. The entire lattice represents a state of no information — denote it by \perp . Likewise $[\text{false}, \text{false}]$ can be identified with *false*, and $[\text{true}, \text{true}]$ with *true*. Then the ordering \leq_t of $\mathcal{K}(L)$ yields operations \wedge and \vee that are those of Kleene's strong three-valued logic, and the negation operation on $\mathcal{K}(L)$ is likewise that of Kleene's logic. The \leq_k ordering puts \perp below both *false* and *true*, which are incomparable. It is the natural notion of approximation in this context.

Kleene's three-valued logic, it was noted above, is a sublogic of *FOUR* — the sublogic of consistent members. In fact this relationship is quite general, and thus the notion of interval approximation yields another approach to bilattices. Proofs of the following can be found in [14].

Theorem 4.3 *Suppose L is a lattice with a de Morgan complement. Then $\mathcal{K}(L)$ is isomorphic to the set of consistent members of an interlaced bilattice \mathcal{B} having a negation and a conflation that commute. If, further, L is distributive (that is, L is a de Morgan lattice), \mathcal{B} will be a distributive bilattice.*

Theorem 4.4 *Suppose \mathcal{B} is a distributive bilattice with a negation and a conflation that commute. Then the set of consistent members of \mathcal{B} is isomorphic to $\mathcal{K}(L)$ where L is a de Morgan lattice.*

Remark As proved in [14], the isomorphisms of the theorems above did not take negations into account. However, it is simple to check that the isomorphisms do respect the negation operations as well.

Thus we have two rather different approaches that both lead to bilattices. In one version we can think of assigning evidence for a sentence as being independent of assigning evidence against it. In the other we work with approximations to the ‘real’ truth values. As the theorems above indicate, these give rise to essentially the same structures. Among these structures are Belnap’s four-valued logic, with Kleene’s three-valued logic as a sublogic. These can be thought of as the setting for much of the semantical work on logic programming that has appeared in the literature, and thus our discussion below, which concerns bilattices in general, applies directly to the usual notions of stable, stationary, and well-founded models.

5 Logic programs

We want to define logic programs in as general a way as possible, so that the results proved later on will be widely applicable. Conventional logic programming is thought of as having $\{false, true\}$ as its intended space of truth values. But since not every query may produce an answer, partial models are often allowed. That is, \perp is added. Likewise, if a mechanism that can produce inconsistent programs is introduced, \top must also be considered. Thus *FOUR* can be thought of as the ‘home’ of ordinary logic programming. In this section we want to extend the notion of logic program, so that a bilattice \mathcal{B} other than *FOUR* can be thought of as the space of truth values. The more general the setting allowed, the more general the results.

Assume \mathcal{B} is a complete (so that quantifiers can be interpreted) bilattice with negation. For starters, we might take a logic program to be a finite set of clauses. As usual in semantic discussions, we will think of a clause as standing for the possibly infinite set of its instances, over a Herbrand base \mathcal{H} . Further, we will take a clause to be something of the form $A \leftarrow B$, where A is the *head* and B is the *body* of the clause. The question is, what should be allowed as head and body.

As a first approximation, we might take a head to be an atomic formula A , and a body to be a finite list L_1, \dots, L_n of atomic formulas and their negations. We can assume equality is available, since we can always add the clause $eq(x, x) \leftarrow$ to a program. Then in the usual way, clauses having the same predicate letter in their heads can be combined, provided we have conjunction, disjunction, and existential quantification explicitly available. For instance, the two clauses:

$$\begin{aligned} A(x) &\leftarrow B(x) \\ A(f(x)) &\leftarrow C(x), D(x) \end{aligned}$$

combine into the single clause:

$$A(y) \leftarrow (\exists x)[(eq(y, x) \wedge B(x)) \vee (eq(y, f(x)) \wedge C(x) \wedge D(x))].$$

Notice also that the explicit quantification makes it possible for all free variables of the clause body to appear in the head, and for the head to contain no function symbols. Since this transformation can always be done, it will simplify things if we generally assume that program bodies can be arbitrary formulas built up using \wedge , \vee , \neg , \exists , that all free variables in clause bodies also occur in clause heads, that clause heads contain no function symbols, and that programs are restricted so that a given predicate letter can occur in the head of only a single clause.

Further it will do no harm, at least semantically, if we also allow \forall , and by doing so we gain a useful advantage. The usual de Morgan laws hold in any bilattice with negation. That is, as might be expected, \wedge and \vee are duals, as are \exists and \forall , so in a clause body we can always push negations all the way inside. Since we also have $\neg\neg x = x$, we can assume the only occurrences of the negation symbol are at the literal level. This will simplify things considerably. There is a significant problem with this, however. Allen Van Gelder has pointed out (private communication) that the definition of well-founded semantics from [30, 31], based on the use of unfounded sets, apparently does not extend to include \forall . The alternating fixpoint approach from [29] does extend to allow it, and it is this approach that we generalize here. It is not the case that all program transformations that preserve meaning when using the original definition of well-founded semantics continue to do so when applied to programs containing \forall , using the alternating fixpoint approach. This *caveat* must be kept in mind in what follows.

Gelfond and Lifschitz, [18], have introduced a kind of classical negation into logic programming. In addition to $\neg P(\mathbf{x})$ they also allow $\sim P(\mathbf{x})$, where \sim is intended to behave classically. In particular, it is allowed in clause heads. In fact, for most purposes $\sim P(\mathbf{x})$ can be treated semantically as if it were a new atom, $P'(\mathbf{x})$. The exception is if a logic program involving this new notion of negation turns out to be inconsistent, in which case they argue that every query should have a yes answer. We feel this may sometimes be undesirable; that contradictions should be localizable. If we have inconsistent information about ducks, it is possible that our information about decimals can still be trusted. To this end we do, in fact, treat $\sim P(\mathbf{x})$ as a new atom, $P'(\mathbf{x})$, but if we want to say that from a contradiction everything follows, we do so by explicitly adding clauses of the form $A \leftarrow B \wedge B'$ to a program, for every atom A and every literal B . But this is not the only solution to the problem of what to do if we get yes answers to both the queries B and B' ? We could, for instance, take B as *overdefined*, and act as if its truth value was \top in the bilattice \mathcal{B} . In other circumstances we might want to take B as *underdefined*, \perp , and act accordingly. To allow for the utmost generality, we admit a mechanism for this into logic programs directly. We allow \otimes and \oplus to appear in clause bodies. Then, if desired, we can include a clause $real_B \leftarrow B \oplus B'$, for instance, to specify explicitly that a yes answer to both B and B' should be thought of as overdefining B . Briefly, rather than determine once and for all how contradictions should be treated, we have added machinery to allow a program writer to specify this. Incidentally, \otimes and \oplus are their own duals under \neg , so we still can assume negations only occur in literals.

In conventional logic programming the truth values *true* and *false* are available. The value *true* is generally implicit; $A \leftarrow$ is treated as if it were $A \leftarrow true$. Likewise having no clause for A is treated as if $A \leftarrow false$ were present. Since we have a more general truth value space in mind, we need a more general mechanism. We want to be able to write something like $A \leftarrow b$ where b is some arbitrary member of the bilattice \mathcal{B} (or maybe a member of a restricted sublogic of \mathcal{B} , such as its exact part). To this end, we simply treat members of \mathcal{B} as atoms, and allow them to appear in clause bodies. Semantically we will, of course, interpret them as designating themselves.

- We call a formula that does not contain a member of \mathcal{B} a *pure* formula.

We now have arrived at a very general notion of logic program, relative to the bilattice \mathcal{B} . It is

broad enough to encompass a variety of generalizations of conventional logic programming. Note, however, that we have not included a mechanism for disjunctive logic programming. This must wait on further research. Now, to summarize precisely.

Definition 5.1

1. A *formula* is an expression built up from literals and members of \mathcal{B} using \wedge , \vee , \otimes , \oplus , \exists and \forall .
2. A *clause* is $P(x_1, \dots, x_n) \leftarrow \varphi(x_1, \dots, x_n)$, where the pure atomic formula $P(x_1, \dots, x_n)$ is the *head*, and the formula $\varphi(x_1, \dots, x_n)$ is the *body*. It is assumed that the free variables of the body are among x_1, \dots, x_n .
3. A *logic program* is a finite set of clauses with no predicate letter appearing in the head of more than one clause.

If \mathcal{P} is a program, by \mathcal{P}^* we mean the set of all ground instances of members of \mathcal{P} , over the Herbrand base.

We will often want to use conventional logic programs as examples. To make this easier, we introduce the following.

Convention A *conventional* logic program is one whose underlying truth value space is the lattice *FOUR*, and which does not involve \otimes , \oplus , \forall , \top or \perp . Such programs can be written in the customary way, using commas to denote conjunction.

6 The Gelfond-Lifschitz transformation

In [6, 17] the *stable model semantics* for logic programs was introduced, and a number of basic results reported. Since this definition and the results will be generalized here, we include the original characterization for reference. For this section only, *program* will mean a conventional logic program. Also the semantics is based on classical logic, *FOUR* restricted to its exact values. In this case, then, we can follow the usual convention of identifying a classical Herbrand model with the set of ground atoms that are true in it.

The semantics of positive programs (those containing no negations in clause bodies) is well understood. In the classic papers [28, 1] it was shown that a positive program has a unique smallest model, and that smallest model is well-behaved and relates nicely to other semantical approaches. The notion of smallest used here is with respect to the subset relation, identifying a model with a set of ground atoms as mentioned above. In order to extend this to programs with negations, Gelfond and Lifschitz proposed the following approach. Start with a set S that is a candidate for a model of program \mathcal{P} ; use S to give meaning to the negative literals of \mathcal{P} , thus converting it into a positive program \mathcal{P}' ; then compute the smallest model of \mathcal{P}' . If this turns out to be S again, S is a good candidate for a model of \mathcal{P} . The conversion of \mathcal{P} into \mathcal{P}' is straightforward. If B is in S we can assume B is true, so any clause in \mathcal{P} containing $\neg B$ in its body is unusable. If B is not in S we can assume $\neg B$ is true, so any occurrence of $\neg B$ in a clause body of \mathcal{P} can be replaced by *true*, or equivalently, dropped. The following more formal treatment is still a sketch, and we adopt the convenience that the program is ground.

Definition 6.1 Let \mathcal{P} be a conventional ground program and let S be a set of ground atoms. The *Gelfond-Lifschitz transformation* of \mathcal{P} relative to S is the positive program \mathcal{P}' arising from \mathcal{P} by:

1. deleting each clause in \mathcal{P} that has a negative literal $\neg B$ in its body, where $B \in S$;
2. from the remaining set of clauses, deleting all negative literals.

Let S' be the least model of \mathcal{P}' . If $S' = S$, then S is called a *stable set*.

Gelfond and Lifschitz show that a stable set, if it exists, must be a model for the program, and so stable sets are often called stable models. In [30, 31], another semantics called the *well-founded semantics* is introduced, and it is shown that if a program has a well-founded (two valued) model, that will be its unique stable model as well. Connections with other semantical approaches were also established, but these will not concern us here. The results just mentioned are generalized later on in this paper.

The semantics sketched above is classical, two-valued. Rather early on, however, the approach was extended to allow partial, or three-valued models. Once this was done, it was possible to show a stable model always exists, though it may be partial. We will consider this point more thoroughly below. For now, we merely remark on an issue of terminology. Stable models were originally intended to be two-valued models. When generalized, three-valued stable models were talked about. Przymusiński introduced the terminology *stationary* for three-valued stable, reserving the term *stable* for the two-valued case. We will not follow this terminology here (preferring to believe the term *stable* has some stability to it). Since we will be extending notions beyond the classical setting, to fairly arbitrary bilattices, we prefer to minimize terminology as far as possible to keep the comprehensible complexity of the paper down.

7 Immediate consequence operators

In [28, 1] an ‘immediate consequence’ operator $T_{\mathcal{P}}$ is associated with each positive, conventional program \mathcal{P} , mapping sets of formulas to sets of formulas. Loosely the idea is that one application of $T_{\mathcal{P}}$ represents carrying out a single step of deduction, based on the program \mathcal{P} . It is shown that $T_{\mathcal{P}}$ is monotonic (in the subset relation), and has a smallest fixed point. That fixed point is taken to be the semantic meaning of the program \mathcal{P} . It is this that was used in the previous section to associate a model with the positive program \mathcal{P}' resulting from the Gelfond-Lifschitz transformation. The operator $T_{\mathcal{P}}$ makes sense for programs \mathcal{P} allowing negation in clause bodies, but then it loses its monotonicity feature, so a fixed point can no longer be guaranteed.

In [7, 8] this approach was extended to a three-valued setting, in which case the operator is generally denoted $\Phi_{\mathcal{P}}$. This operator turns out to be well-behaved, even for programs having negations in clause bodies. But, the intended ordering must be changed. When $T_{\mathcal{P}}$ is used, one thinks of the natural ordering as one in which an increase means more ground atoms are true. When using $\Phi_{\mathcal{P}}$ an increase loosely means more literals acquire classical truth values. Thus the two orderings of a bilattice begin to be visible: $T_{\mathcal{P}}$ uses the truth ordering, $\Phi_{\mathcal{P}}$ uses the knowledge ordering. This will be made more precise below. At any rate, it can be shown that $\Phi_{\mathcal{P}}$ always has a smallest fixed point under the appropriate ordering, whether negations are present or not, and this can be taken as a meaning for the program \mathcal{P} . It is not always a satisfactory meaning, and has led to modifications in [20, 21], and to the alternative approaches sketched earlier.

In [11, 12] the three-valued approach was generalized to the family of bilattices meeting the interlacing conditions. The first thing is to extend the representation of a valuation. In the classical approach, a model is often identified with a set of ground atoms, those that are true in the model. Then the subset relation corresponds to the \leq_t ordering of bilattices. In the three-valued approach, a model is sometimes identified with a (consistent) set of ground literals, where if a ground atom

A is in the set, it means A has the value *true* in the model; if $\neg A$ is in the set, it means A has the value *false* in the model; if neither A nor $\neg A$ is in the set, it means A has the value undefined, or \perp in the model. Under this representation, the subset relation corresponds to the \leq_k ordering of bilattices. But when moving to more general settings, representation by sets of formulas is no longer adequate.

Definition 7.1 Let \mathcal{B} be a bilattice. By a *valuation* in \mathcal{B} we mean a mapping v from pure ground atoms to members of \mathcal{B} . The family of all valuations is denoted $\mathcal{V}(\mathcal{B})$. It is given two pointwise orderings:

1. $v_1 \leq_t v_2$ if and only if $v_1(A) \leq_t v_2(A)$ for every pure ground atom A ;
2. $v_1 \leq_k v_2$ if and only if $v_1(A) \leq_k v_2(A)$ for every pure ground atom A .

The space of valuations itself is a bilattice, and inherits much of the structure of the underlying bilattice. Let us be more precise. First, $\mathcal{V}(\mathcal{B}) = \langle \mathcal{V}(\mathcal{B}), \leq_t, \leq_k \rangle$ is a pre-bilattice. $(v \wedge w)(A) = v(A) \wedge w(A)$, and similarly for the other operations. $\mathcal{V}(\mathcal{B})$ is complete if \mathcal{B} is complete. $\mathcal{V}(\mathcal{B})$ is distributive if \mathcal{B} is distributive. $\mathcal{V}(\mathcal{B})$ is infinitely distributive if \mathcal{B} is infinitely distributive. $\mathcal{V}(\mathcal{B})$ satisfies the interlacing conditions if \mathcal{B} satisfies them. And $\mathcal{V}(\mathcal{B})$ satisfies the infinitary interlacing conditions if \mathcal{B} does. Finally, $\mathcal{V}(\mathcal{B})$ will have a negation or a conflation if \mathcal{B} does. All this is straightforward, and verification is omitted.

The action of a valuation is easily extended from pure ground atoms to all formulas. If $b \in \mathcal{B}$, set $v(b) = b$. Further, set $v(X \wedge Y) = v(X) \wedge v(Y)$, where the \wedge on the right is the meet of \mathcal{B} in the \leq_t ordering. Similarly for \vee , \otimes , \oplus and \neg (assuming \mathcal{B} has a negation operation). Assuming completeness too, we can also set $v((\exists x)\varphi(x)) = \sum\{v(\varphi(t)) \mid \text{all closed terms } t\}$ (the operation is the infinitary join under \leq_t .) Similarly for the universal quantifier.

Connections with earlier work are easy to make. Suppose we use the bilattice *FOUR*, shown in Figure 1. This is infinitely distributive, and has negation and conflation. Hence so does its space of valuations. A valuation in *FOUR* corresponds to a set of pure ground literals in a direct way. Say the valuation v and the set S *correspond* provided:

1. $v(A) = \perp$ if and only if $A \notin S$ and $\neg A \notin S$;
2. $v(A) = \textit{false}$ if and only if $A \notin S$ and $\neg A \in S$;
3. $v(A) = \textit{true}$ if and only if $A \in S$ and $\neg A \notin S$;
4. $v(A) = \top$ if and only if $A \in S$ and $\neg A \in S$.

Then, if the valuation v and the set S correspond, v is exact just in case S contains exactly one of each pure ground atom and its negation; v is consistent just in case S is consistent as a set. This correspondence should be kept in mind in what follows.

Suppose now that \mathcal{B} is a complete bilattice with negation, and \mathcal{P} is a program as in Definition 5.1. We associate with it an operator $\Phi_{\mathcal{P}}$ on the space of valuations. The idea directly extends the three-valued $\Phi_{\mathcal{P}}$ operator mentioned above. Recall that programs are assumed to have at most one defining clause for each predicate letter, and all free variables of a clause body also appear in the clause head. It follows that for a ground atom A , at most one member of \mathcal{P}^* can have A as head (recall that \mathcal{P}^* is the set of ground instances of members of \mathcal{P}).

Definition 7.2 $\Phi_{\mathcal{P}} : \mathcal{V}(\mathcal{B}) \rightarrow \mathcal{V}(\mathcal{B})$ is defined as follows. Let $v \in \mathcal{V}(\mathcal{B})$; $\Phi_{\mathcal{P}}(v)$ is the valuation such that:

1. if the pure ground atom A is not the head of any member of \mathcal{P}^* , $\Phi_{\mathcal{P}}(v)(A) = \text{false}$;
2. if $A \leftarrow B$ occurs in \mathcal{P}^* , $\Phi_{\mathcal{P}}(v)(A) = v(B)$.

This operator naturally generalizes those mentioned above. Suppose \mathcal{B} is the bilattice *FOUR* of Figure 1. Then, if \mathcal{P} is a conventional program without negations, on exact valuations $\Phi_{\mathcal{P}}$ behaves like the operator $T_{\mathcal{P}}$. Further, if \mathcal{P} is a conventional program allowing negations, on consistent valuations $\Phi_{\mathcal{P}}$ behaves like the Kripke-Kleene partial operator of [7].

It can be shown that if \mathcal{B} meets the finitary and infinitary interlacing conditions, $\Phi_{\mathcal{P}}$ is always monotone in the \leq_k ordering of $\mathcal{V}(\mathcal{B})$, and thus must have fixed points. These fixed points are natural generalizations to the bilattice setting of models. Further, $\Phi_{\mathcal{P}}$ will be monotone in the \leq_t ordering if \mathcal{P} has no negations. Relationships between extremal fixed points of $\Phi_{\mathcal{P}}$ are established in [12].

8 The Gelfond-Lifschitz transformation generalized

In Section 7 we saw how the immediate consequence operator $T_{\mathcal{P}}$ generalizes to the bilattice setting. What we want is to produce a similar generalization of the Gelfond-Lifschitz transformation of Section 6. In order to do this it is necessary to shift the point of view from that of program transformations to that of operators on valuations. Suppose we have a conventional logic program \mathcal{P} , and we have a set S of pure ground atoms. Carry out the Gelfond-Lifschitz transformation on \mathcal{P} , relative to S , getting the positive program \mathcal{P}' . This program has a smallest fixed point, S' , which is another set of pure ground atoms. Thus an operator has been associated with the program \mathcal{P} : the operator that turns S into S' . In [29] this was called the *stability* operator. It is the stability operator that we generalize in this section, following [13].

One of the ideas behind the stable model approach, and the well-founded approach too, is to separate the roles of positive and negative information. Thus we extend the immediate consequence operator of Section 7 to reflect this separation. The extended immediate consequence operator will accept two input valuations, one assigning meanings to positive literals, the other to negative literals. This means we will sometimes want to think of a negative literal $\neg A$ as a strangely written atom, with no connections to A . We use the term *\mathcal{B} -pseudo-valuation* for a mapping from pure ground *literals* to \mathcal{B} . When using a pseudo-valuation the value assigned to $\neg A$ can be quite independent of the value assigned to A .

Now we introduce a convenient notation for representing a pseudo-valuation in terms of real valuations.

Definition 8.1 Let v_1 and v_2 be valuations in the bilattice \mathcal{B} with negation. We define a pseudo-valuation in \mathcal{B} , $v_1 \Delta v_2$ as follows. For a pure ground atom A :

$$\begin{aligned} (v_1 \Delta v_2)(A) &= v_1(A) \\ (v_1 \Delta v_2)(\neg A) &= \neg v_2(A) \end{aligned}$$

Pseudo-valuations are extended to non-literals by induction, in the obvious way.

The idea is, in $v_1 \Delta v_2$, v_1 supplies the positive information and v_2 supplies the negative information. Now Definition 7.2 can be suitably generalized. We use Ψ for the new operator, and keep Φ as it was defined earlier, since we will need it as well.

Definition 8.2 The *extended immediate consequence operator*, $\Psi_{\mathcal{P}} : \mathcal{V}(\mathcal{B}) \times \mathcal{V}(\mathcal{B}) \rightarrow \mathcal{V}(\mathcal{B})$, is defined as follows. Let $v_1, v_2 \in \mathcal{V}(\mathcal{B})$; $\Psi_{\mathcal{P}}(v_1, v_2)$ is the valuation such that:

1. if the pure ground atom A is not the head of any member of \mathcal{P}^* , $\Psi_{\mathcal{P}}(v_1, v_2)(A) = \text{false}$;
2. if $A \leftarrow B$ occurs in \mathcal{P}^* , $\Psi_{\mathcal{P}}(v_1, v_2)(A) = (v_1 \Delta v_2)(B)$.

In $\Psi_{\mathcal{P}}$ positive and negative input has been separated. It is easy to see that $\Phi_{\mathcal{P}}(v) = \Psi_{\mathcal{P}}(v, v)$.

Definition 8.3 A mapping f on a partially ordered space is *monotonic* if $x \leq y$ implies $f(x) \leq f(y)$; it is *anti-monotonic* if $x \leq y$ implies $f(x) \geq f(y)$.

If we assume the interlacing conditions hold for \mathcal{B} , under the ordering \leq_k , $\Psi_{\mathcal{P}}$ is monotonic in both arguments. But under the ordering \leq_t , $\Psi_{\mathcal{P}}$ is monotonic in its first argument and anti-monotonic in its second. (Proofs of these and other results are given in the Appendix.) Then under the \leq_t ordering, if we hold the second argument of $\Psi_{\mathcal{P}}$ fixed, we have a monotone operator of its first argument, and this will have a least fixed point. Thus we can make the following definition.

Definition 8.4 The *derived (or stability) operator* of $\Psi_{\mathcal{P}}$ is the single input mapping $\Psi'_{\mathcal{P}}$ given by: $\Psi'_{\mathcal{P}}(v)$ is the smallest fixed point, in the \leq_t ordering, of the mapping $(\lambda x)\Psi_{\mathcal{P}}(x, v)$.

It is the derived operator, $\Psi'_{\mathcal{P}}$, that generalizes the Gelfond-Lifschitz transformation. Suppose \mathcal{P} is a conventional logic program and \mathcal{B} is the bilattice *FOUR* of Figure 1. Let S be a consistent set of ground literals, and let v be the valuation in \mathcal{B} that corresponds to S . Suppose S' is the least fixed point of the program \mathcal{P}' , which is the result of applying the Gelfond-Lifschitz transformation to \mathcal{P} with respect to S . Then the set S' and the valuation $\Psi'_{\mathcal{P}}(v)$ will also correspond. We leave it to you to check this.

Definition 8.5 A *stable valuation* for program \mathcal{P} is a fixed point of $\Psi'_{\mathcal{P}}$.

Note, then, that for the bilattice *FOUR*, a stable valuation that is exact corresponds to a two-valued stable set in the sense of [17], and a stable valuation that is consistent but not necessarily exact corresponds to a stationary set in the sense of Przymusiński.

9 Results

In this section results are presented. Proofs are omitted here, but can be found in the Appendix. An examination of the proofs will show that almost without exception they are simple algebraic consequences of the elementary properties of bilattices, and of the monotonicity or anti-monotonicity of operators. For the rest of this section, \mathcal{B} is a fixed bilattice of truth values, which we assume is infinitely distributive, satisfying the interlacing laws, finitary and infinitary, and with a negation and a conflation that commute. Also \mathcal{P} is a program, meeting the conditions of Definition 5.1. As in the previous section, $\Psi_{\mathcal{P}}$ is the extended immediate consequence operator for program \mathcal{P} , and $\Psi'_{\mathcal{P}}$ is its derived operator.

Theorem 9.1 *The operator $\Psi'_{\mathcal{P}}$ is monotonic in the \leq_k ordering, and anti-monotonic in the \leq_t ordering.*

Since the derived operator is monotonic in the \leq_k direction, and \mathcal{B} is a complete lattice under this ordering, the Knaster-Tarski Theorem ([26]) applies.

Theorem 9.2 *The operator $\Psi'_{\mathcal{P}}$ has a smallest fixed point, denoted $s_{\mathcal{P}}^k$, and a greatest fixed point, denoted $S_{\mathcal{P}}^k$, under the \leq_k ordering.*

This means every program has at least one stable valuation, a fundamental result of [24]. In fact, the least stable valuation $s_{\mathcal{P}}^k$ has been singled out before as playing a special role. We will say more about it below. The greatest stable valuation, $S_{\mathcal{P}}^k$, has been ignored, essentially because consistency was required of logic programs. Once various degrees of inconsistency are allowed, the essential symmetry of the situation becomes apparent. There is both a least *and a greatest* stable valuation, in the \leq_k sense.

Although stable valuations exist, we have not yet established in what sense they are models for the program. We give a definition of model that is, in fact, rather strong.

Definition 9.3 A valuation v is a *model* for the program \mathcal{P} provided, for each pure ground atom A , if $A \leftarrow B$ in \mathcal{P}^* , $v(A) = v(B)$, and otherwise $v(A) = \text{false}$.

Now the role of the immediate consequence operator, $\Phi_{\mathcal{P}}$, is clear. Obviously v is a model of \mathcal{P} if and only if v is a fixed point of $\Phi_{\mathcal{P}}$. Then the result of [17], that stable sets are models, has a full generalization.

Theorem 9.4 *Every stable valuation is a model. That is, every fixed point of $\Psi'_{\mathcal{P}}$ is also a fixed point of $\Phi_{\mathcal{P}}$.*

The behavior of $\Psi'_{\mathcal{P}}$ with respect to the \leq_k ordering is thus rather straightforward. With respect to the \leq_t ordering, however, things are more complicated, since the operator is anti-monotonic in this direction. As it happens, there is a straightforward modification of the Knaster-Tarski theorem that deals with this. It seems to have first appeared in [32], where it was used as the basis of an interesting Theory of Truth. It lies behind the alternating fixpoint approach to the well-founded model, as presented in [29]. It was also used in [2] to show the existence of stable classes.

Theorem 9.5 *Suppose f is anti-monotonic on the complete lattice L . Then there are two members of L , μ and ν , such that:*

1. μ and ν are the least and the greatest fixed points of f^2 ;
2. f oscillates between μ and ν in the sense that $f(\mu) = \nu$ and $f(\nu) = \mu$;
3. if x and y are also points of L between which f oscillates, x and y lie between μ and ν .

Now, $\Psi'_{\mathcal{P}}$ is anti-monotonic under \leq_t , so by the theorem above, it has two extreme oscillation points in this ordering.

Definition 9.6 The extreme oscillation points of $\Psi'_{\mathcal{P}}$ are denoted $s_{\mathcal{P}}^t$ and $S_{\mathcal{P}}^t$, with $s_{\mathcal{P}}^t \leq_t S_{\mathcal{P}}^t$.

The valuations $s_{\mathcal{P}}^t$ and $S_{\mathcal{P}}^t$ need not themselves be stable. But two stable valuations are easily derivable from them.

Theorem 9.7 *The valuations $s_{\mathcal{P}}^t \otimes S_{\mathcal{P}}^t$ and $s_{\mathcal{P}}^t \oplus S_{\mathcal{P}}^t$ are both fixed points of $\Psi'_{\mathcal{P}}$, and hence are stable models.*

The model $s_{\mathcal{P}}^t \otimes S_{\mathcal{P}}^t$ plays a special role in several ways. For conventional programs, using the bilattice *FOUR*, it is the *well-founded model* [30, 31], following the construction of [29]. If we think of $s_{\mathcal{P}}^t$ and $S_{\mathcal{P}}^t$ as under and over estimates for a model, the well-founded model, $s_{\mathcal{P}}^t \otimes S_{\mathcal{P}}^t$, is the most the estimates agree on — their consensus. The model $s_{\mathcal{P}}^t \oplus S_{\mathcal{P}}^t$ has not been considered before

in the literature, primarily because inconsistent truth values have been ruled out. It is, however, a natural dual to the well-founded model, accepting everything either of the extreme oscillation points of Ψ'_p have to contribute.

An important connection was established between well-founded semantics and stable semantics in [24, 22], for conventional programs. Essentially it says that the well-founded model is the ‘simplest’ stable model, where the ordering used is essentially \leq_k of *FOUR*. That result not only generalizes to bilattices, but dualizes as well.

Theorem 9.8 *The extremal fixed points of Ψ'_p under \leq_k are related to the extremal oscillation points under \leq_t as follows:*

1. $s_p^k = s_p^t \otimes S_p^t$;
2. $S_p^k = s_p^t \oplus S_p^t$.

Thus the distribution of the family of stable models is neatly bounded. All stable models are between s_p^k and S_p^k in the \leq_k ordering, with these as least and greatest stable models in this ordering. All stable models are between s_p^t and S_p^t in the \leq_t ordering, with these points themselves not necessarily included. This is summarized in Figure 2.

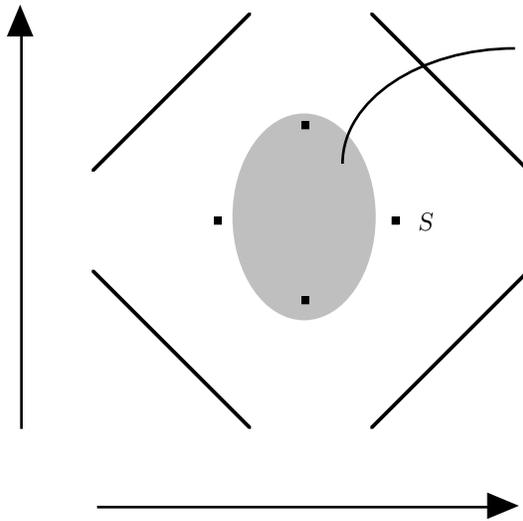


Figure 2: Stable models in the space of valuations

Rather remarkably, the result presented in Theorem 9.8 has a counterpart, relating the extremal points of the stability operator the other way around.

Theorem 9.9

1. $s_p^t = s_p^k \wedge S_p^k$;
2. $S_p^t = s_p^k \vee S_p^k$.

The two Theorems 9.8 and 9.9 should be compared with Theorem 7.7 of [12]. Finally, it was shown in [30, 31] that if a conventional program has a two-valued (classical) well-founded model, it is also the unique stable model. The analog for bilattices of a classical model is a model that is exact in the space of valuations. Then, with a restriction on programs added, the result extends directly.

Definition 9.10 We call a program \mathcal{P} *consistent* if no program body contains the operator \oplus , and the only members of \mathcal{B} that occur in program bodies are consistent ones.

Conventional logic programs are obviously consistent ones. Part of the justification for the terminology is contained in the following.

Theorem 9.11 *If \mathcal{P} is a consistent program, $s_{\mathcal{P}}^t$ and $S_{\mathcal{P}}^t$ are consistent in the space of valuations.*

Now the Gelfond, Lifschitz result extends to the following.

Corollary 9.12 *Suppose \mathcal{P} is a consistent program, and the (well-founded) model $s_{\mathcal{P}}^k = s_{\mathcal{P}}^t \otimes S_{\mathcal{P}}^t$ is exact. Then it is the unique stable model.*

10 Examples

We give a few small examples to illustrate the results of the previous section. The first two concern the bilattice *FOUR* and conventional logic programs.

Example 10.1 Using the bilattice *FOUR*, consider the logic program:

$$\begin{aligned} A &\leftarrow \neg B \\ B &\leftarrow \neg A \end{aligned}$$

This has four stable models. Two of them are $s_{\mathcal{P}}^k$ and $S_{\mathcal{P}}^k$, the least and greatest in the \leq_k ordering. The model $s_{\mathcal{P}}^k$ is the well-founded model. Neither of these is exact. In addition there are two more stable models, both exact (we denote them v and w). Finally, $s_{\mathcal{P}}^t$ and $S_{\mathcal{P}}^t$ exist as valuations, though neither is a model. These are given in the following table.

	$s_{\mathcal{P}}^t$	$S_{\mathcal{P}}^t$	$s_{\mathcal{P}}^k$	$S_{\mathcal{P}}^k$	v	w
A	false	true	\perp	\top	true	false
B	false	true	\perp	\top	false	true

The relative positions of these valuations in the space of all valuations is shown in Figure 3. Note that v and w are incomparable in both orderings.

Example 10.2 Again we use the bilattice *FOUR*. The program is:

$$\begin{aligned} B &\leftarrow \neg A \\ C &\leftarrow \neg B \vee D \\ D &\leftarrow \neg D \end{aligned}$$

For this program, the extremal valuations are shown in the following table.

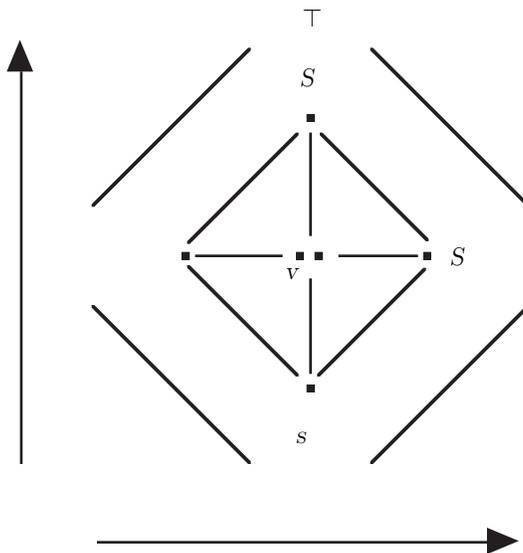


Figure 3: Distribution of valuations

	$s_{\mathcal{P}}^t$	$S_{\mathcal{P}}^t$	$s_{\mathcal{P}}^k$	$S_{\mathcal{P}}^k$
A	<i>false</i>	<i>false</i>	<i>false</i>	<i>false</i>
B	<i>true</i>	<i>true</i>	<i>true</i>	<i>true</i>
C	<i>false</i>	<i>true</i>	\perp	\top
D	<i>false</i>	<i>true</i>	\perp	\top

The stable model $S_{\mathcal{P}}^k$ is not consistent in the space of valuations. Even so, it does not degenerate to triviality. Under it, C and D are overdefined, but A and B still have useful values. Allowing inconsistencies does not have to mean surrendering to chaos.

Example 10.3 For this rather different example we use a bilattice based on the unit interval: take $\mathcal{B} = L \odot L$ where L is $[0, 1]$ with the usual ordering. For this bilattice additional operations besides the usual ones can be introduced, to some profit. We are interested in what we will call *attenuation operators*. Specifically, for each $r \in [0, 1]$ we introduce a mapping $\mathcal{A}_r : \mathcal{B} \times \mathcal{B} \rightarrow \mathcal{B}$ where $\mathcal{A}_r(\langle x, y \rangle) = \langle rx, 1 - r(1 - y) \rangle$. It is simple to check that each attenuation operator maps exact values to exact values, consistent values to consistent values, and is monotonic under both \leq_t and \leq_k .

When writing programs we can allow the syntactic use of \mathcal{A}_r , treating it the same way we do negation. We extend valuations to attenuation operators in the obvious way: $v(\mathcal{A}_r\varphi) = \mathcal{A}_r(v(\varphi))$. Now, assume the Herbrand base consists of $\{a, f(a), f(f(a)), \dots\}$ and consider the following program, \mathcal{P} .

$$P(f(x)) \leftarrow \mathcal{A}_{0.9}(P(x))$$

In [27] van Emden introduced a logic programming language with the unit interval as its set of truth values. What we are considering now is a generalization of that. Numerical factors were

associated with each clause, in the van Emden paper. These correspond to the attenuation operators introduced above. More precisely, if we only apply attenuation operators to entire clause bodies, only allow exact members of \mathcal{B} in clause bodies, and restrict logical operators to \wedge , \vee and \exists , the system is essentially that of van Emden.

The program \mathcal{P} has an attenuation operator in a clause body, something that was not mentioned in Section 9. However, we observed above that attenuation operators are monotone in both bilattice orderings. It is a simple task to check that, as a consequence, almost every proof in the Appendix continues to be valid even if attenuation operators are present. The exceptions are Theorem 9.11 and Corollary 9.12 which use the notion of a consistent program, a notion that needs re-defining when attenuation operators are present. Thus the entire machinery of stable and well-founded models applies. In fact, the program \mathcal{P} is rather well-behaved. A direct calculation shows that the smallest and greatest stable models in the \leq_k ordering coincide, and are exact. Hence the well-founded model is the unique stable model. It is not, however, the only model. The least and greatest fixed points of $\Phi_{\mathcal{P}}$ are different. Along with the well-founded model, they are partly given in the following table, under the names $\mu_{\mathcal{P}}$ and $\nu_{\mathcal{P}}$.

	$s_{\mathcal{P}}^k = S_{\mathcal{P}}^k$	$\mu_{\mathcal{P}}$	$\nu_{\mathcal{P}}$
$P(a)$	$\langle 0, 1 \rangle$	$\langle 0, 0 \rangle$	$\langle 1, 1 \rangle$
$P(f(a))$	$\langle 0.9, 0.1 \rangle$	$\langle 0, 0.1 \rangle$	$\langle 0.9, 1 \rangle$
$P(f^2(a))$	$\langle 0.09, 0.91 \rangle$	$\langle 0.09, 0.1 \rangle$	$\langle 0.9, 0.91 \rangle$
$P(f^3(a))$	$\langle 0.819, 0.181 \rangle$	$\langle 0.09, 0.181 \rangle$	$\langle 0.819, 0.91 \rangle$
$P(f^4(a))$	$\langle 0.1629, 0.8371 \rangle$	$\langle 0.1629, 0.181 \rangle$	$\langle 0.819, 0.8371 \rangle$

Notice that the largest (non-stable) model, $\nu_{\mathcal{P}}$, is not consistent in the space of valuations. Still, it displays considerable structure. The values it assigns to ground atoms are uniformly not consistent in \mathcal{B} , but as n increases the value assigned to $P(f^n(a))$ grows ‘less inconsistent.’ Incidentally, in all cases, as n increases the value assigned to $P(f^n(a))$ approaches $\langle 0.4736\dots, 0.5263\dots \rangle$ as a limit. This value is exact.

11 Conclusion

Stable models, and in particular the well-founded model, have received much attention in the literature recently. There is a general feeling that this approach comes closer to capturing the meaning of logic programs than other attempts. We have seen that the family of stable models has some remarkably simple bounds, bounds that are closely inter-related. This adds to their interest. But there is still much to be done. For instance, how does this work extend to disjunctive logic programs? In addition there are general structural questions: we have established bounds on the family of stable models; is there more structure to be discovered? The fixed points of a monotone operator in a complete lattice form another complete lattice. The derived operator $\Psi'_{\mathcal{P}}$ is monotonic under \leq_k and its fixed points are the stable models. What additional information is given by knowing the stable models constitute a complete lattice themselves?

We have broadened the usual context of logic programming semantics by insisting that inconsistent truth values be allowed, and be taken seriously. We feel this is fundamental. Inconsistencies in information are facts of life; we often believe that different inconsistencies really are different. It should be possible to model this feeling in the semantics. We hope others will follow up on this point.

Finally, we have worked with the general family of bilattices, rather than confining things to the setting of conventional logic programming. This adds no complexity to the proofs in the Appendix. Instead it helps make clear the rather simple assumptions on which the stable model semantics rests. Among the variety of bilattices, there are many of intrinsic interest. We would like to see program languages developed and experimented with that use other bilattices, in the sense that conventional logic programming uses *FOUR*. There may be useful things here.

Appendix — Proofs

Proofs of the results presented in Section 9 are collected together here. It should quickly become apparent that these are all algebraic consequences of elementary properties of bilattices and the stability operator. We begin by assuming enough properties of the underlying bilattice so that things are well-behaved.

Assumption For this section, \mathcal{B} is an infinitely distributive bilattice, satisfying the interlacing laws, finitary and infinitary, and with a negation and a conflation that commute.

Lemma *Let $v_1, v_2, w_1, w_2 \in \mathcal{V}(\mathcal{B})$ and let A be a ground formula with all negations at the literal level. Then:*

1. *if $v_1 \leq_k v_2$ then $v_1(A) \leq_k v_2(A)$;*
2. *if A does not contain \neg and if $v_1 \leq_t v_2$ then $v_1(A) \leq_t v_2(A)$;*
3. *if $v_1 \leq_k v_2$ and $w_1 \leq_k w_2$ then $v_1 \Delta w_1 \leq_k v_2 \Delta w_2$;*
4. *if $v_1 \leq_t v_2$ and $w_2 \leq_t w_1$ then $v_1 \Delta w_1 \leq_t v_2 \Delta w_2$.*

Proof The Lemma holds if A is atomic by definition of the orderings in $\mathcal{V}(\mathcal{B})$, Definition 7.1. It extends to literals using the definition of negation, Definition 3.5, and Definition 8.1. Then the extension to more complex formulas is by a straightforward induction on the degree of A , using the interlacing conditions, Definition 3.4. ■

Proposition *In the space of valuations:*

1. *the operator $\Phi_{\mathcal{P}}$ (Definition 7.2) is monotonic under \leq_k ;*
2. *the operator $\Psi_{\mathcal{P}}$ (Definition 8.2) is monotonic in both arguments under \leq_k , and under the ordering \leq_t it is monotonic in its first argument and anti-monotonic in its second.*

Proof Suppose $v_1 \leq_k v_2$ and let A be a pure ground atom; we want to show $\Phi_{\mathcal{P}}(v_1)(A) \leq_k \Phi_{\mathcal{P}}(v_2)(A)$. If A does not occur as the head of any member of \mathcal{P}^* , trivially $\Phi_{\mathcal{P}}(v_1)(A) = \Phi_{\mathcal{P}}(v_2)(A)$, since both are *false*, so the result is immediate. Otherwise, if $A \leftarrow B \in \mathcal{P}^*$, $\Phi_{\mathcal{P}}(v_1)(A) = v_1(B)$, and similarly for v_2 , so we need that $v_1(B) \leq_k v_2(B)$, and this is the case by part 1 of the previous Lemma. This establishes Item 1. Item 2 is by a similar argument, using parts 3 and 4 of the Lemma. ■

This Proposition is central for virtually all of what follows. Now, before we continue we take a moment to recall the two usual proofs of the Knaster-Tarski Theorem.

Suppose L is a complete lattice and f is monotone on L . The Knaster-Tarski Theorem says, among other things, that f has a smallest fixed point, s . There are two ways of ‘constructing’ this fixed point, and each leads to a technique for proving things about it.

In one approach, the least fixed point of f is shown to be $\bigwedge\{x \in L \mid f(x) \leq x\}$. It follows that if $f(x) \leq x$, then $s \leq x$. This is a convenient way of showing upper bounds on s , and we use it in the proof immediately following.

In the other standard approach one produces a (generally transfinite) sequence of members of L as follows. f_0 is set to be the least member of L . For an ordinal α , $f_{\alpha+1}$ is set to be $f(f_\alpha)$. And for a limit ordinal λ , f_λ is set to be $\bigvee_{\alpha < \lambda} f_\alpha$. This sequence is increasing. The limit of the sequence, $\bigvee_\alpha f_\alpha$, is the least fixed point of f , and the fixed point is attained at some first ordinal, the *closure ordinal* of f . This yields another method of proof: by transfinite induction. If it can be shown that each member of the f_α sequence has some property, then the least fixed point, s , also has that property. We use this technique also in the proof below.

Theorem 9.1 *The operator Ψ'_P is anti-monotonic in the \leq_t ordering, but is monotonic in the \leq_k ordering.*

Proof Suppose $v_1 \leq_t v_2$. By the previous Proposition, Ψ_P is anti-monotonic in its second argument, so

$$\Psi_P(\Psi'_P(v_1), v_2) \leq_t \Psi_P(\Psi'_P(v_1), v_1).$$

Since $\Psi'_P(v_1)$ is a fixed point of $(\lambda x)\Psi_P(x, v_1)$, this yields:

$$\Psi_P(\Psi'_P(v_1), v_2) \leq_t \Psi'_P(v_1).$$

Since $\Psi'_P(v_2)$ is the *least* fixed point of $(\lambda x)\Psi_P(x, v_2)$ under \leq_t , it follows from one of the proofs of the Knaster-Tarski Theorem that

$$\Psi'_P(v_2) \leq_t \Psi'_P(v_1).$$

This establishes the first part of the Theorem.

Now suppose $v_1 \leq_k v_2$. We will show $\Psi'_P(v_1) \leq_k \Psi'_P(v_2)$. To this end we define two transfinite sequences of valuations, a_α and b_α , as follows. $a_0 = b_0$ is the always *false* valuation, the least in the \leq_t ordering. $a_{\alpha+1} = \Psi_P(a_\alpha, v_1)$ and $b_{\alpha+1} = \Psi_P(b_\alpha, v_2)$. Finally, for a limit ordinal λ , $a_\lambda = \bigvee_{\alpha < \lambda} a_\alpha$ and $b_\lambda = \bigvee_{\alpha < \lambda} b_\alpha$. Both sequences are increasing in the \leq_t ordering since Ψ'_P is monotone in its first argument. The a_α sequence has $\Psi'_P(v_1)$ as its limit, while the b_α sequence has $\Psi'_P(v_2)$ as its limit, so to prove the second assertion of the theorem it is enough to establish that $a_\alpha \leq_k b_\alpha$ for every ordinal α .

If $\alpha = 0$, a_α and b_α are the same; this case is trivial.

Suppose $a_\alpha \leq_k b_\alpha$. Then $a_{\alpha+1} = \Psi_P(a_\alpha, v_1) \leq_k \Psi_P(b_\alpha, v_2) = b_{\alpha+1}$, using the monotonicity of Ψ_P in both arguments, under \leq_k .

Finally, suppose $a_\alpha \leq_k b_\alpha$ for every $\alpha < \lambda$. It follows that $\bigvee_{\alpha < \lambda} a_\alpha \leq_k \bigvee_{\alpha < \lambda} b_\alpha$, using the fact that $\mathcal{V}(\mathcal{B})$ satisfies the infinitary interlacing conditions. ■

With this theorem shown, the existence of a smallest fixed point s_P^k , and a biggest fixed point S_P^k , for Ψ'_P , under the \leq_k ordering, follows immediately from the Knaster-Tarski Theorem. Thus we have established

Theorem 9.2 *The operator Ψ'_P has a smallest fixed point, denoted s_P^k , and a greatest fixed point, denoted S_P^k , under the \leq_k ordering.*

The next result is now shown by a simple calculation.

Theorem 9.4 *Every stable valuation is a model. That is, every fixed point of $\Psi'_{\mathcal{P}}$ is also a fixed point of $\Phi_{\mathcal{P}}$.*

Proof Suppose s is a fixed point of $\Psi'_{\mathcal{P}}$. Then

$$\Phi_{\mathcal{P}}(s) = \Psi_{\mathcal{P}}(s, s) = \Psi_{\mathcal{P}}(\Psi'_{\mathcal{P}}(s), s) = \Psi'_{\mathcal{P}}(s) = s.$$

■

We omit the proof of Theorem 9.5, which essentially amounts to applying the Knaster-Tarski argument to the monotonic operator f^2 . Using the Theorem, however, the existence of the two extreme oscillation points $s_{\mathcal{P}}^t$ and $S_{\mathcal{P}}^t$ follows immediately. Before continuing with results from Section 9, we need a Proposition, and a few Lemmas leading up to it.

Lemma *In a distributive bilattice, $x = (x \wedge \perp) \oplus (x \vee \perp)$, and also $false \otimes true = \perp$.*

Proof

$$\begin{aligned} (x \wedge \perp) \oplus (x \vee \perp) &= [x \oplus (x \vee \perp)] \wedge [\perp \oplus (x \vee \perp)] \\ &= [(x \oplus x) \vee (x \oplus \perp)] \wedge [(\perp \oplus x) \vee (\perp \oplus \perp)] \\ &= [x \vee x] \wedge [x \vee \perp] \\ &= x \wedge (x \vee \perp) \\ &= x \end{aligned}$$

For the second part, $false \leq_t \perp$ since $false$ is the smallest member of the bilattice under the \leq_t ordering. Then using the interlacing conditions, which hold in any distributive bilattice,

$$false \otimes true \leq_t \perp \otimes true = \perp.$$

Similarly $\perp \leq_t true$, so

$$\perp = false \otimes \perp \leq_t false \otimes true.$$

■

Incidentally, there are three more equations like $false \otimes true = \perp$ that hold in distributive bilattices: $false \oplus true = \top$; $\perp \wedge \top = false$; and $\perp \vee \top = true$. These have similar proofs, but we do not need them here.

Lemma *In a distributive bilattice, if $a \leq_t b \leq_t c$ then:*

1. $(a \wedge \perp) \otimes (c \vee \perp) \leq_k \perp$;
2. $(a \vee \perp) \otimes c \leq_k b$;
3. $(a \wedge \perp) \otimes (c \wedge \perp) \leq_k b$.

Proof Since $\perp \leq_k false$, by the interlacing conditions, $a \wedge \perp \leq_k a \wedge false = false$. Similarly $c \vee \perp \leq_k true$. Then, again by the interlacing conditions, $(a \wedge \perp) \otimes (c \vee \perp) \leq_k false \otimes true$. In any distributive bilattice, $false \otimes true = \perp$. This establishes part 1.

Next, using the hypothesis, and interlacing, $(a \vee \perp) \otimes c \leq_k (a \vee b) \otimes c = b \otimes c \leq_k b$.

Finally, $(a \wedge \perp) \otimes (c \wedge \perp) \leq_k (a \wedge \perp) \otimes (c \wedge b) = (a \wedge \perp) \otimes b \leq_k b$. ■

Finally, the result we chiefly need.

Proposition *In a distributive bilattice, if $a \leq_t b \leq_t c$ then $a \otimes c \leq_k b$.*

Proof Using the Lemmas above, and interlacing,

$$\begin{aligned}
a \otimes c &= [(a \wedge \perp) \oplus (a \vee \perp)] \otimes c \\
&= [(a \wedge \perp) \otimes c] \oplus [(a \vee \perp) \otimes c] \\
&\leq_k [(a \wedge \perp) \otimes c] \oplus b \\
&= [(a \wedge \perp) \otimes ((c \wedge \perp) \oplus (c \vee \perp))] \oplus b \\
&= [(a \wedge \perp) \otimes (c \wedge \perp)] \oplus [(a \wedge \perp) \otimes (c \vee \perp)] \oplus b \\
&\leq_k b \oplus \perp \oplus b \\
&= b
\end{aligned}$$

■

Remark The Proposition above is really one of four related results. Using a similar proof, and appropriate versions of the Lemmas preceding, the following can also be shown.

- if $a \leq_t b \leq_t c$ then $b \leq_k a \oplus c$;
- if $a \leq_k b \leq_k c$ then $a \wedge c \leq_t b$;
- if $a \leq_k b \leq_k c$ then $b \leq_t a \vee c$.

We will need these items later on. Now to return to results from Section 9.

Theorem 9.7 *The valuations $s_{\mathcal{P}}^t \otimes S_{\mathcal{P}}^t$ and $s_{\mathcal{P}}^t \oplus S_{\mathcal{P}}^t$ are both fixed points of $\Psi'_{\mathcal{P}}$, and hence are stable models.*

Proof Using monotonicity of $\Psi'_{\mathcal{P}}$ in the \leq_k ordering,

$$\Psi'_{\mathcal{P}}(s_{\mathcal{P}}^t \otimes S_{\mathcal{P}}^t) \leq_k \Psi'_{\mathcal{P}}(s_{\mathcal{P}}^t) = S_{\mathcal{P}}^t$$

$$\Psi'_{\mathcal{P}}(s_{\mathcal{P}}^t \otimes S_{\mathcal{P}}^t) \leq_k \Psi'_{\mathcal{P}}(S_{\mathcal{P}}^t) = s_{\mathcal{P}}^t$$

and it follows that

$$\Psi'_{\mathcal{P}}(s_{\mathcal{P}}^t \otimes S_{\mathcal{P}}^t) \leq_k s_{\mathcal{P}}^t \otimes S_{\mathcal{P}}^t.$$

Also, $s_{\mathcal{P}}^t \leq_t S_{\mathcal{P}}^t$, so by interlacing,

$$s_{\mathcal{P}}^t = s_{\mathcal{P}}^t \otimes s_{\mathcal{P}}^t \leq_t s_{\mathcal{P}}^t \otimes S_{\mathcal{P}}^t \leq_t S_{\mathcal{P}}^t \otimes S_{\mathcal{P}}^t = S_{\mathcal{P}}^t$$

and then by the anti-monotonicity of $\Psi'_{\mathcal{P}}$ in the \leq_t ordering,

$$\Psi'_{\mathcal{P}}(S_{\mathcal{P}}^t) \leq_t \Psi'_{\mathcal{P}}(s_{\mathcal{P}}^t \otimes S_{\mathcal{P}}^t) \leq_t \Psi'_{\mathcal{P}}(s_{\mathcal{P}}^t)$$

and so

$$s_{\mathcal{P}}^t \leq_t \Psi'_{\mathcal{P}}(s_{\mathcal{P}}^t \otimes S_{\mathcal{P}}^t) \leq_t S_{\mathcal{P}}^t.$$

Now by the previous Proposition,

$$s_{\mathcal{P}}^t \otimes S_{\mathcal{P}}^t \leq_k \Psi'_{\mathcal{P}}(s_{\mathcal{P}}^t \otimes S_{\mathcal{P}}^t).$$

This proves half the Theorem; the other half has a dual proof. ■

Theorem 9.8 *The extremal fixed points of $\Psi'_{\mathcal{P}}$ under \leq_k can be calculated from the extremal oscillation points under \leq_t as follows:*

1. $s_{\mathcal{P}}^k = s_{\mathcal{P}}^t \otimes S_{\mathcal{P}}^t$;
2. $S_{\mathcal{P}}^k = s_{\mathcal{P}}^t \oplus S_{\mathcal{P}}^t$.

Proof The valuation $s_{\mathcal{P}}^k$ is the least fixed point of $\Psi'_{\mathcal{P}}$ in the \leq_k ordering. Theorem 9.7 showed that $s_{\mathcal{P}}^t \otimes S_{\mathcal{P}}^t$ is a fixed point, hence

$$s_{\mathcal{P}}^k \leq_k s_{\mathcal{P}}^t \otimes S_{\mathcal{P}}^t.$$

In the other direction, $s_{\mathcal{P}}^t$ and $S_{\mathcal{P}}^t$ are extremal oscillation points in the \leq_t ordering. But $s_{\mathcal{P}}^k$, being a fixed point, is trivially an oscillation point. Hence

$$s_{\mathcal{P}}^t \leq_t s_{\mathcal{P}}^k \leq_t S_{\mathcal{P}}^t$$

and it follows from the Proposition earlier that

$$s_{\mathcal{P}}^t \otimes S_{\mathcal{P}}^t \leq_k s_{\mathcal{P}}^k.$$

This is half the theorem. Once again, the other half is dual. ■

Next, a result about monotone mappings in general, before we return to Section 9.

Proposition *If f is a monotone mapping on the complete lattice L , then f and f^2 have the same least and greatest fixed points.*

Proof We show the result for least fixed points; the other half has a dual proof. Let a be the least fixed point of f , and let b be the least fixed point of f^2 .

Since every fixed point of f is also a fixed point of f^2 , $f^2(a) = a$. Since b is the least fixed point of f^2 , $b \leq a$.

If x is a fixed point of f^2 , so is $f(x)$, because $f^2(f(x)) = f(f^2(x)) = f(x)$. Then $f(b)$ is a fixed point of f^2 , and since b is the least fixed point of f^2 , $b \leq f(b)$. By monotonicity, $f(b) \leq f^2(b)$, or $f(b) \leq b$. Since a is the least fixed point of f , it follows that $a \leq b$. ■

Now we continue with results from Section 9.

Theorem 9.9

1. $s_{\mathcal{P}}^t = s_{\mathcal{P}}^k \wedge S_{\mathcal{P}}^k$;
2. $S_{\mathcal{P}}^t = s_{\mathcal{P}}^k \vee S_{\mathcal{P}}^k$.

Proof The mapping $\Psi'_{\mathcal{P}}$ is monotone in the \leq_k ordering, with $s_{\mathcal{P}}^k$ and $S_{\mathcal{P}}^k$ as its least and greatest fixed points in this ordering. Trivially, $(\Psi'_{\mathcal{P}})^2$ is also monotone in this ordering, and by the Proposition above it has the same least and greatest fixed points. But $(\Psi'_{\mathcal{P}})^2$ is also monotone in the \leq_t ordering, with $s_{\mathcal{P}}^t$ and $S_{\mathcal{P}}^t$ as its least and greatest fixed points in this ordering. Now,

$$s_{\mathcal{P}}^k \wedge S_{\mathcal{P}}^k \leq_t s_{\mathcal{P}}^k$$

hence

$$(\Psi'_{\mathcal{P}})^2(s_{\mathcal{P}}^k \wedge S_{\mathcal{P}}^k) \leq_t (\Psi'_{\mathcal{P}})^2(s_{\mathcal{P}}^k) = s_{\mathcal{P}}^k$$

and similarly

$$(\Psi'_{\mathcal{P}})^2(s_{\mathcal{P}}^k \wedge S_{\mathcal{P}}^k) \leq_t S_{\mathcal{P}}^k$$

so

$$(\Psi'_{\mathcal{P}})^2(s_{\mathcal{P}}^k \wedge S_{\mathcal{P}}^k) \leq_t s_{\mathcal{P}}^k \wedge S_{\mathcal{P}}^k.$$

Then since $s_{\mathcal{P}}^t$ is the least fixed point of $(\Psi'_{\mathcal{P}})^2$ in the \leq_t ordering,

$$s_{\mathcal{P}}^t \leq_t s_{\mathcal{P}}^k \wedge S_{\mathcal{P}}^k.$$

Further, since $s_{\mathcal{P}}^t$ is a fixed point of $(\Psi'_{\mathcal{P}})^2$, and $s_{\mathcal{P}}^k$ and $S_{\mathcal{P}}^k$ are the least and greatest fixed points of this operator in the \leq_k ordering,

$$s_{\mathcal{P}}^k \leq_k s_{\mathcal{P}}^t \leq_k S_{\mathcal{P}}^k$$

and it follows by an earlier Remark that

$$s_{\mathcal{P}}^k \wedge S_{\mathcal{P}}^k \leq_t s_{\mathcal{P}}^t.$$

■

For the remaining results, recall that a member x of a bilattice with conflation is *consistent* if $x \leq_k -x$, and a consistent program is one in which program bodies do not contain \oplus , and contain only consistent members of \mathcal{B} .

Theorem 9.11 *If \mathcal{P} is a consistent program, $s_{\mathcal{P}}^t$ and $S_{\mathcal{P}}^t$ are consistent in the space of valuations.*

Proof The first thing we need is that if \mathcal{P} is a consistent program, $\Psi'_{\mathcal{P}}$ maps consistent valuations to consistent valuations. We begin with some background.

Consistent valuations are those that are consistent in the bilattice of valuations. Equivalently, they are the valuations that map ground atoms to consistent members of \mathcal{B} . It is shown in [12] (Proposition 3.6) that the consistent members of a bilattice meeting the conditions we are assuming contain *false* and are closed under negation and both finitary and infinitary versions of \wedge , \vee , and \otimes . More can be said than this, but it is all we need here.

It follows from what was just said that if v is a consistent valuation and B is a ground formula built up from atoms and consistent members of \mathcal{B} , using \wedge , \vee , \neg , \exists , \forall , and \otimes , then $v(B)$ will be a consistent member of \mathcal{B} . It follows that if \mathcal{P} is a consistent program, and v and w are consistent valuations, then $\Psi_{\mathcal{P}}(v, w)$ is a consistent valuation.

Next, assume w is a consistent valuation. $\Psi'_{\mathcal{P}}(w)$ is the least fixed point, in the \leq_t ordering, of the monotone operator $(\lambda x)\Psi_{\mathcal{P}}(x, w)$. One approximates to this least fixed point via a transfinite sequence of steps. The initial member, the always *false* valuation, is consistent. If stage α yields a consistent valuation, so will stage $\alpha + 1$, by the previous paragraph. And at limit ordinals we use the sup operation, which in this case is the infinitary version of \vee , and the collection of consistent members is closed under this operation. It follows that every member of the transfinite sequence is consistent, and hence $\Psi'_{\mathcal{P}}(w)$ is consistent. Thus $\Psi'_{\mathcal{P}}$ does map consistent valuations to consistent valuations.

The valuation $s_{\mathcal{P}}^t$ is the least fixed point, in the \leq_t ordering, of the monotone operator $(\Psi'_{\mathcal{P}})^2$. By the preceding paragraph, if \mathcal{P} is a consistent program, $(\Psi'_{\mathcal{P}})^2$ maps consistent valuations to consistent valuations. Then again, every member of the transfinite sequence of approximations to the least fixed point of it will be consistent, hence so will $s_{\mathcal{P}}^t$. The result about $S_{\mathcal{P}}^t$ can be shown by a dual argument. ■

Corollary 9.12 *Suppose \mathcal{P} is a consistent program, and the well-founded model $s_{\mathcal{P}}^k$ is exact. Then it is the unique stable model.*

Proof If \mathcal{P} is a consistent program, $s_{\mathcal{P}}^t$ is consistent, and so $s_{\mathcal{P}}^t \leq_k -s_{\mathcal{P}}^t$. If $s_{\mathcal{P}}^k$ is exact, $s_{\mathcal{P}}^k = -s_{\mathcal{P}}^k$. Now, $s_{\mathcal{P}}^k \leq_k s_{\mathcal{P}}^t$ since $s_{\mathcal{P}}^k = s_{\mathcal{P}}^t \otimes S_{\mathcal{P}}^t$. Then $-s_{\mathcal{P}}^t \leq_k -s_{\mathcal{P}}^k$ and it follows, under the present assumptions, that $s_{\mathcal{P}}^t \leq_k s_{\mathcal{P}}^k$. Hence $s_{\mathcal{P}}^k = s_{\mathcal{P}}^t$.

In a similar way, $s_{\mathcal{P}}^k = S_{\mathcal{P}}^t$. Then also, $S_{\mathcal{P}}^k = s_{\mathcal{P}}^t \oplus S_{\mathcal{P}}^t = s_{\mathcal{P}}^k \oplus s_{\mathcal{P}}^k = s_{\mathcal{P}}^k$. Since all stable models lie between $s_{\mathcal{P}}^k$ and $S_{\mathcal{P}}^k$, there is only one stable model. ■

References

- [1] APT, K. R., AND VAN EMDEN, M. Contributions to the theory of logic programming. *Journal of the Assoc. for Comp. Mach.* 29 (1982), 841–862.
- [2] BARAL, C., AND SUBRAHMANIAN, V. S. Dualities between alternative semantics for logic programming and nonmonotonic reasoning. In *Logic Programming and Non-monotonic Reasoning*, A. Nerode, W. Marek, and V. S. Subrahmanian, Eds. MIT Press, 1991, pp. 69–86.
- [3] BELNAP, JR., N. D. A useful four-valued logic. In *Modern Uses of Multiple-Valued Logic*, J. M. Dunn and G. Epstein, Eds. D. Reidel, 1977.
- [4] DUNN, J. M. Intuitive semantics for first-degree entailments and coupled trees. *Philosophical Studies* 29 (1976), 149–168.
- [5] DUNN, J. M. Relevance logic and entailment. In *Handbook of Philosophical Logic, Volume III*, D. Gabbay and F. Guenther, Eds. D. Reidel, 1986, ch. 3, pp. 117–224.
- [6] FINE, K. The justification of negation as failure. In *Logic, Methodology and Philosophy of Science VIII* (Amsterdam, 1989), J. E. Fenstad, I. T. Frolov, and R. Hilpinen, Eds., North-Holland, pp. 263–301.
- [7] FITTING, M. C. A Kripke/Kleene semantics for logic programs. *Journal of Logic Programming* 2 (1985), 295–312.
- [8] FITTING, M. C. Partial models and logic programming. *Theoretical Computer Science* 48 (1987), 229–255.
- [9] FITTING, M. C. Bilattices and the theory of truth. *Journal of Philosophical Logic* 18 (1989), 225–256.
- [10] FITTING, M. C. Negation as refutation. In *Proceedings of the Fourth Annual Symposium on Logic in Computer Science* (1989), R. Parikh, Ed., IEEE, pp. 63–70.
- [11] FITTING, M. C. Bilattices in logic programming. In *The Twentieth International Symposium on Multiple-Valued Logic* (1990), G. Epstein, Ed., IEEE, pp. 238–246.
- [12] FITTING, M. C. Bilattices and the semantics of logic programming. *Journal of Logic Programming* 11 (1991), 91–116.
- [13] FITTING, M. C. Well-founded semantics, generalized. In *Logic Programming, Proceedings of the 1991 International Symposium* (Cambridge, MA, 1991), V. Saraswat and K. Ueda, Eds., MIT Press, pp. 71–84.

- [14] FITTING, M. C. Kleene's logic, generalized. *Journal of Logic and Computation* 1 (1992), 797–810.
- [15] FITTING, M. C. Many-valued modal logics, II. *Fundamenta Informaticae* 17 (1992), 55–73.
- [16] FITTING, M. C. Kleene's three-valued logics and their children. In *Proceedings of the Bulgarian Kleene '90 Conference* (Forthcoming).
- [17] GELFOND, M., AND LIFSCHITZ, V. The stable model semantics for logic programming. In *Proc. of the Fifth Logic Programming Symposium* (Cambridge, MA, 1988), R. Kowalski and K. Bowen, Eds., MIT Press, pp. 1070–1080.
- [18] GELFOND, M., AND LIFSCHITZ, V. Logic programs with classical negation. In *Proc. of the Seventh International Conference on Logic Programming* (Cambridge, MA, 1990), D. H. D. Warren and P. Szeredi, Eds., MIT Press, pp. 579–597.
- [19] GINSBERG, M. L. Multivalued logics: a uniform approach to reasoning in artificial intelligence. *Computational Intelligence* 4 (1988), 265–316.
- [20] KUNEN, K. Negation in logic programming. *Journal of Logic Programming* 4 (1987), 289–308.
- [21] KUNEN, K. Some remarks on the completed database. In *Logic Programming, Proc. of the Fifth Intl. Conf. and Symp.* (1988), R. A. Kowalski and K. A. Bowen, Eds., The MIT Press, pp. 978–992.
- [22] PRZYMUSINSKI, T. C. Extended stable semantics for normal and disjunctive programs. In *Proc. of the 7th Int. Conf. on Logic Programming* (1990), D. H. D. Warren and P. Szeredi, Eds., MIT Press, pp. 459–477.
- [23] PRZYMUSINSKI, T. C. Stationary semantics for disjunctive logic programs and deductive databases. In *Logic Programming, Proc. of the 1990 North American Conf.* (1990), S. Debray and M. Hermenegildo, Eds., The MIT Press, pp. 40–59.
- [24] PRZYMUSINSKI, T. C. Well-founded semantics coincides with three-valued stable-semantics. *Fundamenta Informaticae* 13 (1990), 445–463.
- [25] RASIOWA, H., AND MAREK, W. On reaching consensus by groups of intelligent agents. In *Methodologies for Intelligent Systems, 4* (1989), Z. W. Ras, Ed., North-Holland, pp. 234–243.
- [26] TARSKI, A. A lattice-theoretical theorem and its applications. *Pacific Journal of Mathematics* 5 (1955), 285–309.
- [27] VAN EMDEN, M. Quantitative deduction and its fixpoint theory. *Journal of Logic Programming* 3 (1986), 37–53.
- [28] VAN EMDEN, M., AND KOWALSKI, R. The semantics of predicate logic as a programming language. *Journal of the Assoc. for Comp. Mach.* 23 (1976), 733–742.
- [29] VAN GELDER, A. The alternating fixpoint of logic programs with negation. In *Proc. 8th ACM Symp. on Principles of Database Systems* (Philadelphia, 1989), ACM, pp. 1–10.
- [30] VAN GELDER, A., ROSS, K. A., AND SCHLIPF, J. S. Unfounded sets and well-founded semantics for general logic programs. In *Proc. Seventh Symp. on Principles of Database Systems* (1988), pp. 221–230.

- [31] VAN GELDER, A., ROSS, K. A., AND SCHLIPF, J. S. The well-founded semantics for general logic programs. *JACM* 38 (1991), 620–650.
- [32] YABLO, S. Truth and reflection. *Journal of Philosophical Logic* 14 (1985), 297–349.