CLOSURE OPERATORS, FRAMES AND
NEATEST REPRESENTATIONS

ROB EGROT

Abstract

Given a poset P and a standard closure operator $\Gamma : \wp(P) \to \wp(P)$, we give a necessary and sufficient condition for the lattice of Γ-closed sets of $\wp(P)$ to be a frame in terms of the recursive construction of the Γ-closure of sets. We use this condition to show that, given a set $\mathcal{U}$ of distinguished joins from P, the lattice of $\mathcal{U}$-ideals of P fails to be a frame if and only if it fails to be σ-distributive, with σ depending on the cardinalities of sets in $\mathcal{U}$. From this we deduce that if a poset has the property that whenever $a \wedge (b \vee c)$ is defined for $a, b, c \in P$ it is necessarily equal to $(a \wedge b) \vee (a \wedge c)$, then it has an $(\omega, 3)$-representation.

1. Introduction

Schein [15] defines a meet-semilattice S to be distributive if it satisfies the first-order definable condition that whenever $a \wedge (b \vee c)$ is defined, $(a \wedge b) \vee (a \wedge c)$ is also defined and the two are equal. A dual definition can, of course, be made for join-semilattices. Here, for simplicity, all semilattices are considered to be meet-semilattices, and the order duals to our results are left unstated. Note that Schein's version of distributivity is strictly weaker than the notion of distributivity introduced by Grätzer and Schmidt [10] (see, for example, [20]), although they coincide for lattices. Schein's distributivity (hereafter referred to as 3-distributivity, for reasons that will become clear later) is equivalent to the property of being embeddable into a powerset algebra via a semilattice homomorphism preserving all existing binary joins [2].

Schein's 3-distributivity can be generalised to the concept of α-distributivity for cardinals α (see Definition 3.1). This notion has been studied when $\alpha = n < \omega$ [11], when $\alpha = \omega$ [4, 20] and when α is any regular cardinal [12]. Note that if $m, n \leq \omega$ with $m < n$, then n-distributivity trivially implies m-distributivity, but the converse is not true [13]. A recurring theme in these investigations is that a semilattice S is α-distributive if and only if the complete lattice of downsets of S that are closed

under existing α-small joins is a frame (that is, it satisfies the complete distributivity condition that $x \wedge \bigvee Y = \bigvee_Y (x \wedge y)$ for all elements x and subsets Y).

Notions relating to distributivity have also been studied in the more general setting of partially ordered sets (posets) [12, 18]. Note that the concepts in [18] and [12] are not equivalent, even for semilattices (that of [18] is a straightforward generalisation of the distributivity of Grätzer and Schmidt [10], while that of [12] is more closely related to the approach of Schein [15]).

As the transition from lattices to semilattices causes previously equivalent formulations of distributivity to diverge, so too does the transition from semilattices to posets. The following example is of particular interest to us. For a semilattice, being an ω-distributive poset, as defined in [12], is equivalent to being an ω-distributive semilattice in the sense used here [12, Proposition 2.3], which is, in turn, equivalent to being embeddable into a powerset algebra via a map preserving all existing finite meets and joins [2]. However, this is not true for arbitrary posets, as the ω-distributivity of [12] is strictly stronger than the powerset algebra embedding property [12, Example 4.2 and Theorem 4.8].

The property of being embeddable into a powerset algebra via an embedding preserving meets and joins of certain cardinalities has been studied using the terminology *representable* [3, 7, 14] (see Definition 4.1). A notable departure from the semilattice case is that the first-order theory of representable posets is considerably more complex. For example, while the class of (m, n)-representable posets is elementary for all m, n with $2 < m, n \leq \omega$ [7], explicit first-order axioms are not known and the class cannot be finitely axiomatised [6]. This contrasts with the semilattice case where intuitive first-order axioms are known and only a finite number are required to ensure (ω, n)-representability for finite n [2].

In [3], a condition for posets, which we will refer to as **MD**, generalising the distributive property for semilattices, is defined (see Definition 4.2). The authors conjecture [3, Section 3] that this condition is sufficient to ensure that a poset has a representation preserving existing binary joins and all existing finite meets (they call this a *neatest representation*). It is easily seen that this is not a *necessary* condition for such a representation (see, for example, [7, Example 1.5]).

In the original phrasing of [3, Definition 1], it is unclear whether a neatest representation must preserve *arbitrary* existing meets or only finite ones. From the context, we assume the latter, as with the former definition the conjecture is false. To see this, note that every Boolean algebra satisfies **MD** and, indeed, the stronger condition that if $x \wedge \bigvee Y$ is defined, then $\bigvee_Y (x \wedge y)$ is also defined and the two are equal (a Boolean algebra also satisfies the dual to this condition, so a complete Boolean algebra is structurally both a frame and a co-frame). However, it is well known that a Boolean algebra has a representation preserving arbitrary meets and/or joins if and only if it is atomic [1, Corollary 1].

However, if we only demand that *finite* meets are preserved, then the conjecture is indeed correct, which we prove in this note. The main step in our proof is a result classifying the standard closure operators on $\wp(P)$ whose lattices of closed sets

are frames as being precisely those that can be constructed using a certain recursive procedure (Theorem 3.5). This can be viewed as a generalisation of [12, Theorem 2.7]. The proof of the conjecture about neatest representations is then an easy corollary.

We note that the argument in the solution of this conjecture boils down to a straightforward extension of the forward direction of the argument for semilattices in [11, Theorem 2] to the poset setting. Nevertheless, the result is, perhaps, surprising given that the class of posets with this kind of representation cannot be finitely axiomatised (see Section 5 for a discussion of this).

We proceed as follows. In Section 2, we introduce join-specifications and the concept of the radius of a standard closure operator and assorted supporting terminology. In Section 3, we use this to prove the central theorem (Theorem 3.5) and, in the fourth section, we prove a general version of the conjecture on neatest representations (Corollary 4.6). The final section is devoted to a short discussion of the implications of this for the theory of poset representations and some indirectly related questions in complexity theory.

2. Join-specifications and standard closure operators

We begin with a brief introduction to closure operators and their relationship with poset completions. This topic has been studied in detail and a recent survey can be found in [8]. Nevertheless, it will be useful to present some well known results with a consistent terminology and to make explicit some facts that are only implicit elsewhere. We will give specific references when appropriate.

DEFINITION 2.1 (Closure operator). Given a set X, a closure operator on X is a map Γ from $\wp(X)$ to itself that is extensive, monotone and idempotent. That is, for all $S, T \subseteq X$:

(1) $S \subseteq \Gamma(S)$;
(2) $S \subseteq T \implies \Gamma(S) \subseteq \Gamma(T)$; and
(3) $\Gamma(\Gamma(S)) = \Gamma(S)$.

Given a set $S \subseteq P$, we write $S^{\downarrow} = \{p \in P : p \leq s \text{ for some } s \in S\}$, and $p^{\downarrow}$ is a shorthand for $\{p\}^{\downarrow}$. Given a poset P, we are interested in closure operators on P. In particular, we are interested in *standard* closure operators, that is, closure operators such that $\Gamma(\{p\}) = p^{\downarrow}$ for all $p \in P$.

DEFINITION 2.2 (Join-completion). Given a poset P, a join-completion of P is a complete lattice L and an order embedding $e : P \to L$ such that $e[P]$ (the image of P under e) is join-dense in L.

Given a standard closure operator $\Gamma : \wp(P) \to \wp(P)$, the Γ-closed sets form a complete lattice when ordered by inclusion (which we denote by $\Gamma[\wp(P)]$), and the map $\phi_{\Gamma} : P \to \Gamma[\wp(P)]$ defined by $\phi_{\Gamma}(p) = p^{\downarrow}$ is a join-completion of P (we usually omit the subscript and just write ϕ). If I is an indexing set and C_i is a Γ-closed set for all $i \in I$, then $\bigwedge_I C_i = \bigcap_I C_i$ and $\bigvee_I C_i = \Gamma(\bigcup_I C_i)$ (see, for example,

[8, Section 2.1]). Conversely, given a join-completion $e : P \to L$, the sets $e^{-1}[x^{\downarrow}]$ for $x \in L$ define the closed sets of a standard closure operator $\Gamma_e : \wp(P) \to \wp(P)$. This well known connection can be expressed as the following proposition.

PROPOSITION 2.3. *There is a dual isomorphism between the complete lattice of standard closure operators on P (ordered by pointwise inclusion) and the complete lattice of ($e[P]$ preserving isomorphism classes of) join-completions of P (ordered by defining $(e_1 : P \to L_1) \le (e_2 : P \to L_2)$ if and only if there is an order embedding $\psi : L_1 \to L_2$ such that $\psi \circ e_1 = e_2$; if ψ exists it will necessarily preserve all existing meets).*

PROOF. A direct argument is straightforward. See, for example, [8, Propositions 2.1 and 2.19], or the introduction to [18], for equivalent results. □

We are interested in join-completions $e : P \to L$ where the embedding e preserves certain existing joins from P. Often this is done by making some uniform selection, such as the joins of all sets smaller than some fixed cardinal (when they exist; see [16]). We intend to be more general and, to that end, we make the following definition.

DEFINITION 2.4 (Join-specification). Let P be a poset. Let $\mathcal{U}$ be a subset of $\wp(P)$. Then $\mathcal{U}$ is a join-specification (of P) if it satisfies the following conditions:

(1) $\bigvee S$ exists in P for all $S \in \mathcal{U}$;
(2) $\{p\} \in \mathcal{U}$ for all $p \in P$, and
(3) $\emptyset \notin \mathcal{U}$.

Definition 2.4 is similar to that of a *subset selection* (see [8, Section 2]). The difference is that we demand that the selection contains the singletons and that every selected set has a defined join. This serves to tidy up some of the later definitions.

DEFINITION 2.5 (Radius of $\mathcal{U}$). Given a join-specification $\mathcal{U}$, we define the radius of $\mathcal{U}$ to be the smallest cardinal σ such that $\sigma > |S|$ for all $S \in \mathcal{U}$.

DEFINITION 2.6 ($\Gamma_{\mathcal{U}}$). Given a join-specification $\mathcal{U}$ with radius σ and $S \subseteq P$, we define the following subsets of P using transfinite recursion.

• $\Gamma_0(S) = S^{\downarrow}$.
• If $\alpha + 1$ is a successor ordinal, then $\Gamma_{\alpha+1}(S) = \{\bigvee T : T \in \mathcal{U} \text{ and } T \subseteq \Gamma_\alpha(S)\}^{\downarrow}$.
• If λ is a limit ordinal, $\Gamma_\lambda(S) = \bigcup_{\beta < \lambda} \Gamma_\beta(S)$.

We define $\Gamma_{\mathcal{U}} : \wp(P) \to \wp(P)$ by $\Gamma_{\mathcal{U}}(S) = \Gamma_\chi(S)$ for all $S \in \wp(P)$, where χ is the smallest regular cardinal with $\sigma \le \chi$.

DEFINITION 2.7 ($\mathcal{U}$-ideal). Given a join-specification $\mathcal{U}$, we define a $\mathcal{U}$-ideal of P to be a downset closed under joins from $\mathcal{U}$. We define the empty set to be a $\mathcal{U}$-ideal.

The following is a generalisation of [12, Proposition 1.2].

PROPOSITION 2.8. *If $\mathcal{U}$ is a join-specification, then $\Gamma_{\mathcal{U}}$ is the standard closure operator taking $S \subseteq P$ to the smallest $\mathcal{U}$-ideal containing S.*

PROOF. Let $S \in \wp(P) \setminus \{\emptyset\}$ and let I be the smallest $\mathcal{U}$-ideal containing S. Then it must be that $\Gamma_{\mathcal{U}}(S) \subseteq I$ by the closure requirements of I. It is easy to see that $\Gamma_{\mathcal{U}}(S)$ is a downset, so it remains only to show that $\Gamma_{\mathcal{U}}(S)$ is closed under joins from $\mathcal{U}$. So let $X \in \mathcal{U}$ and suppose that $X \subseteq \Gamma_{\mathcal{U}}(S)$. By definition of $\Gamma_{\mathcal{U}}$ we have $X \subseteq \bigcup_{\beta < \chi} \Gamma_{\beta}(S)$, so for each $x \in X$, there is some $\beta_x < \chi$ with $x \in \Gamma_{\beta_x}(S)$. Since χ is regular, there must be β' with $\beta_x \leq \beta' < \chi$ for all $x \in X$, and so $\bigvee X \in \Gamma_{\beta'+1}(S) \subseteq \Gamma_{\chi}(S) = \Gamma_{\mathcal{U}}(S)$, as required. It is straightforward to show that the function taking sets to the smallest $\mathcal{U}$-ideal containing them is a standard closure operator. □

PROPOSITION 2.9. *Given a join-specification $\mathcal{U}$, the canonical map $\phi : P \to \Gamma_{\mathcal{U}}[\wp(P)]$ preserves arbitrary existing meets and the joins of all sets from $\mathcal{U}$.*

PROOF. This is well known, but we give a short proof for the sake of completeness. First, recall that arbitrary intersections of Γ-closed sets are also Γ-closed. So if $\bigwedge S = t$ in P, then $\bigwedge \phi[S] = \bigcap_S s^{\downarrow} = t^{\downarrow} = \phi(t)$, as required. Preservation of joins from $\mathcal{U}$ follows from Lemma 2.11 and Corollary 2.13(2) below. □

Note that ϕ may also preserve joins of sets that are not in $\mathcal{U}$. Given a join-specification $\mathcal{U}$, there will generally be more than one join-completion preserving the specified joins. We are interested in $\Gamma_{\mathcal{U}}[\wp(P)]$, which is, in fact, the largest such join-completion ([17] attributes this result to [5]). This is easily seen by noting that if $e : P \to L$ is a join-completion such that $e(\bigvee S) = \bigvee e[S]$ for some $S \subseteq P$, then whenever $x \in L$ and $S \subseteq e^{-1}[x^{\downarrow}]$ it must be that $\bigvee S \in e^{-1}[x^{\downarrow}]$ as, otherwise, $\bigvee e[S] \neq e(\bigvee S)$. So, in particular, if $e : P \to L$ preserves joins from $\mathcal{U}$, then $e^{-1}[x^{\downarrow}]$ is a $\mathcal{U}$-ideal for all $x \in L$. Since L is isomorphic to $\{e^{-1}[x^{\downarrow}] : x \in L\}$ considered as a lattice ordered by inclusion, the result follows, as we can think of L as being a subset of the set of all $\mathcal{U}$-ideals.

DEFINITION 2.10 ($\mathcal{U}_{\Gamma}$). Any standard closure operator $\Gamma : \wp(P) \to \wp(P)$ defines a join-specification $\mathcal{U}_{\Gamma}$ by

$$S \in \mathcal{U}_{\Gamma} \iff \bigvee S \text{ exists, and } S \subseteq C \implies \bigvee S \in C \text{ for all } \Gamma\text{-closed sets } C.$$

LEMMA 2.11. *For $\mathcal{U}_{\Gamma}$ as in Definition 2.10 and the canonical $\phi : P \to \Gamma[\wp(P)]$,*

$$S \in \mathcal{U}_{\Gamma} \iff \bigvee S \text{ exists and } \phi\left(\bigvee S\right) = \bigvee \phi[S].$$

PROOF. Let $S \in \mathcal{U}_{\Gamma}$. Then $\bigvee \phi[S]$ is the smallest Γ-closed set containing S, so must therefore contain $\bigvee S$ and is, indeed, equal to $(\bigvee S)^{\downarrow} = \phi(\bigvee S)$. Conversely, if $\bigvee S$ exists and $\bigvee \phi[S] = \phi(\bigvee S)$, then every Γ-closed set containing S must contain $\bigvee S$, and so $S \in \mathcal{U}_{\Gamma}$, by definition. □

The join-specifications of a poset P are a subset of $\wp(\wp(P))$, and as this subset is closed under taking arbitrary unions and intersections, they form a complete lattice when ordered by inclusion. This leads us to the following result.

PROPOSITION 2.12. *Let $\mathcal{J}$ be the lattice of join-specifications of P and let C be the lattice of standard closure operators on P. Define $f : \mathcal{J} \to C$ and $g : C \to \mathcal{J}$ by $f(\mathcal{U}) = \Gamma_{\mathcal{U}}$ and $g(\Gamma) = \mathcal{U}_\Gamma$. Then f and g form a Galois connection between C and $\mathcal{J}$. That is, for all $\Gamma \in C$ and $\mathcal{U} \in \mathcal{J}$, $\Gamma_{\mathcal{U}}(T) \leq \Gamma(T)$ for all $T \in \wp(P) \iff \mathcal{U} \subseteq \mathcal{U}_\Gamma$.*

PROOF. Let $\Gamma \in C$ and $\mathcal{U} \in \mathcal{J}$. Suppose that $\Gamma_{\mathcal{U}} \leq \Gamma$. Since $\mathcal{J}$ is ordered by inclusion we let $S \in \mathcal{U}$, and we aim to show that $S \in \mathcal{U}_\Gamma$. Now $\Gamma_{\mathcal{U}}(S)$ is the smallest $\mathcal{U}$-ideal containing S, so $\bigvee S \in \Gamma_{\mathcal{U}}(S)$, and since $\Gamma_{\mathcal{U}}(S) \leq \Gamma(S)$, this means that any Γ-closed set containing S must contain $\bigvee S$. But then, by definition, $S \in \mathcal{U}_\Gamma$, which is what we wished to show.

For the converse, suppose that $\mathcal{U} \subseteq \mathcal{U}_\Gamma$ and let $T \in \wp(P)$. We aim to show that $\Gamma_{\mathcal{U}}(T) \subseteq \Gamma(T)$. Since $\Gamma_{\mathcal{U}}(T)$ is the smallest $\mathcal{U}$-ideal containing T, it is sufficient to show that $\Gamma(T)$ is also a $\mathcal{U}$-ideal. But since $\mathcal{U} \subseteq \mathcal{U}_\Gamma$, this follows directly from the definition of $\mathcal{U}_\Gamma$. □

COROLLARY 2.13.

(1) *For all standard closure operators $\Gamma : \wp(P) \to \wp(P)$ we have $\Gamma_{\mathcal{U}_\Gamma} \leq \Gamma$, but it is not necessarily true that $\Gamma = \Gamma_{\mathcal{U}_\Gamma}$.*

(2) *For all join-specifications $\mathcal{U}$ of P we have $\mathcal{U} \subseteq \mathcal{U}_{\Gamma_{\mathcal{U}}}$, but it is not necessarily true that $\mathcal{U} = \mathcal{U}_{\Gamma_{\mathcal{U}}}$.*

(3) *If $\mathcal{U}' = \mathcal{U}_{\Gamma_{\mathcal{U}}}$, then $\Gamma_{\mathcal{U}} = \Gamma_{\mathcal{U}'}$.*

(4) *If $\Gamma' = \Gamma_{\mathcal{U}_\Gamma}$, then $\mathcal{U}_\Gamma = \mathcal{U}_{\Gamma'}$.*

PROOF. This all follows from the fact that f and g from Proposition 2.12 form a Galois connection. Examples 2.14 and 2.15 below witness lack of equality for parts (1) and (2), respectively. □

EXAMPLE 2.14. Let P be the three element antichain $\{a, b, c\}$ and let the Γ-closed sets be $\emptyset, \{a\}, \{b\}, \{c\}$ and $\{a, b, c\}$. Then $\mathcal{U}_\Gamma = \{\{a\}, \{b\}, \{c\}\}$ (so $\Gamma[\wp(P)]$ is the MacNeille completion of P) and the set of $\mathcal{U}_\Gamma$-ideals of P is just $\wp(P)$ (so $\Gamma_{\mathcal{U}_\Gamma}[\wp(P)]$ is the Alexandroff completion composed of all downsets of P, in this case). Then, for example, $\Gamma(\{a, b\}) = \{a, b, c\}$, but $\Gamma_{\mathcal{U}_\Gamma}(\{a, b\}) = \{a, b\}$.

Example 2.14 also demonstrates that not every standard closure operator arises from a join-specification. This is because the only join-specification on P is $\{\{a\}, \{b\}, \{c\}\}$ and, as we saw in the example, the induced closure operator produces the Alexandroff completion and not, for example, the MacNeille completion. This issue is also discussed in [17, Section 2].

EXAMPLE 2.15. Let P be a poset containing elements x, x', y, y' and z. Let the nontrivial orderings be $x < x' < z$ and $y < y' < z$, so that $x \vee y = x' \vee y' = z$. Let $\mathcal{U} = \{\{x\}, \{x'\}, \{y\}, \{y'\}, \{z\}, \{x, y\}\}$ be a join-specification. Then $\{x', y'\} \in \mathcal{U}_{\Gamma_{\mathcal{U}}} \setminus \mathcal{U}$.

It follows from Corollary 2.13 that different join-specifications can define the same closure operator. Given a standard closure operator Γ arising from a join-specification, while there is not necessarily a minimal generating join-specification, Definition 2.5

allows us to define a class of join-specifications generating Γ whose sets have the smallest possible maximum size. This puts an upper bound on the number of iterations required in a recursive construction of Γ (although we cannot expect to do better than ω, even if all the sets in the generating join-specification have bounded finite size). These minimal join-specifications (Definition 2.17 below) are also relevant when discussing the distributivity of the lattice of Γ-closed sets (see Corollary 3.6).

DEFINITION 2.16 (Radius of Γ). Given a standard closure operator Γ such that $\Gamma = \Gamma_{\mathcal{U}}$ for some join-specification $\mathcal{U}$, we define the radius of Γ to be the minimum of the set $\{\chi : \chi \text{ is the radius of a join-specification } \mathcal{U}' \text{ of } P \text{ with } \Gamma_{\mathcal{U}'} = \Gamma\}$.

DEFINITION 2.17 (Minimal join-specification). A join-specification $\mathcal{U}$ of P is minimal if the radius of $\mathcal{U}$ is equal to the radius of $\Gamma_{\mathcal{U}}$.

The following technical lemma will be used in the next section.

LEMMA 2.18. *Let $\mathcal{U}$ be a join-specification. Then the following hold for all $S \in \wp(P)$.*

(1)   *If $\Gamma_{\mathcal{U}}(S) = p^{\downarrow}$, then $p = \bigvee S$.*
(2)   *If $p = \bigvee S$ and $p \in \Gamma_{\mathcal{U}}(S)$, then $S \in \mathcal{U}_{\Gamma_{\mathcal{U}}}$.*

PROOF. For the first part, note that p must be an upper bound for S. If $p \neq \bigvee S$, then S has another upper bound q with $p \not\leq q$. But $\Gamma_{\mathcal{U}}(S) \subseteq p^{\downarrow} \cap q^{\downarrow}$ by Proposition 2.8, and thus $p \notin \Gamma_{\mathcal{U}}(S)$, which is a contradiction. For the second part, note that since $\bigvee S = p$, Definition 2.10 gives $S \in \mathcal{U}_{\Gamma_{\mathcal{U}}}$ if and only if $S \subseteq C \implies p \in C$ for all $\Gamma_{\mathcal{U}}$-closed sets C. Since $\Gamma_{\mathcal{U}}(S)$ is the smallest $\Gamma_{\mathcal{U}}$-ideal containing S, if $p \in \Gamma_{\mathcal{U}}(S)$, then every $\Gamma_{\mathcal{U}}$-ideal containing S must also contain p and the proof is complete.   □

3. When is a lattice of $\mathcal{U}$-ideals a frame?

As mentioned in the introduction, definitions of distributivity in semilattices modelled on that of Schein [15] give rise to results that can be stated in our terminology as *a semilattice S is α-distributive if and only if the lattice of $\mathcal{U}_{\alpha}$-ideals of S is α-distributive*, where $\mathcal{U}_{\alpha}$ contains all sets smaller than α whose joins are defined (where α is some finite or regular cardinal) [4, 11, 12]. Moreover, it turns out that if this lattice of $\mathcal{U}_{\alpha}$-ideals is α-distributive, then it will be a frame (see Definition 3.1 below). We can extend this to posets and arbitrary join-specifications, in the sense that if the lattice of $\mathcal{U}$-ideals of P fails to be a frame, then it must be because distributivity fails for the embedded images of some element of P and some set in $\mathcal{U}$ (see Corollary 3.7). The key result is Theorem 3.5, which can be seen as a partial generalisation of [12, Theorem 2.7] to arbitrary join-specifications.

DEFINITION 3.1 (α-distributive). Given a cardinal α, we say that a lattice (or a semilattice) L is α-distributive if given $\{x\} \cup Y \subseteq L$ such that $|Y| < \alpha$, if $x \wedge \bigvee Y$ exists, then $\bigvee_{Y}(x \wedge y)$ also exists and the two are equal. Note that when $3 \leq n \leq \omega$, in the lattice case, n-distributivity is just distributivity. When L is a complete lattice and L is α-distributive for all α, we say that L is a *frame*.

From now on we fix a minimal join-specification $\mathcal{V}$ and we define $\Phi = \Gamma_{\mathcal{V}}$. Let σ be the radius of $\mathcal{V}$ and let χ be the smallest regular cardinal with $\sigma \leq \chi$. Similarly, let σ' be the radius of $\mathcal{U}_\Phi$ (recall Definition 2.10) and let χ' be the smallest regular cardinal with $\sigma' \leq \chi'$.

DEFINITION 3.2 (Υ). Given $S \subseteq P$, we define the following subsets of P by transfinite recursion.

- $\Upsilon_0(S) = S^{\downarrow}$.
- If $\alpha + 1$ is a successor ordinal, $\Upsilon_{\alpha+1}(S) = \{\bigvee T : T \in \mathcal{U}_\Phi \text{ and } T \subseteq \Upsilon_\alpha(S)\}$.
- If λ is a limit ordinal, $\Upsilon_\lambda(S) = \bigcup_{\beta < \lambda} \Upsilon_\beta(S)$.

We define $\Upsilon : \wp(P) \to \wp(P)$ by $\Upsilon(S) = \Upsilon_{\chi'}(S)$ for all $S \in \wp(P)$.

This differs from Definition 2.6 in that we do not close downwards during the successor steps. Note the use of χ' in place of χ which is important in the proof of Corollary 3.6. Note that Υ will not necessarily be a closure operator as it may not be idempotent.

LEMMA 3.3. $\Upsilon(S) \subseteq \Phi(S)$ for all $S \in \wp(P)$.

PROOF. We have $\Phi = \Gamma_{\mathcal{V}} = \Gamma_{\mathcal{U}_\Phi}$ by Proposition 2.13(3) and, by definition of Υ, it follows that $\Upsilon(S) \subseteq \Gamma_{\mathcal{U}_\Phi}(S)$ for all $S \in \wp(P)$.   □

LEMMA 3.4. Let J be an indexing set, and let $I, K_j \in \Phi[\wp(P)]$ for all $j \in J$. Then $I \cap \Upsilon(\bigcup_J K_j) \subseteq \Phi(\bigcup_J (I \cap K_j))$.

PROOF. We note that $\Upsilon(\bigcup_J K_j) = \bigcup_{\alpha < \chi'} \Upsilon_\alpha(\bigcup_J K_j)$. We proceed by showing that $I \cap \Upsilon_\alpha(\bigcup_J K_j) \subseteq \Phi(\bigcup_J (I \cap K_j))$ for all α using transfinite induction on α. If $\alpha = 0$, the result is trivial so consider the successor ordinal $\alpha + 1$ and assume the appropriate inductive hypothesis. Let $p \in I \cap \Upsilon_{\alpha+1}(\bigcup_J K_j)$. Then $p = \bigvee T$ for some $T \in \mathcal{U}_\Phi$ with $T \subseteq \Upsilon_\alpha(\bigcup_J K_j)$ and, since $p \in I$, it follows that $T \subseteq I \cap \Upsilon_\alpha(\bigcup_J K_j)$. By the inductive hypothesis this means that $T \subseteq \Phi(\bigcup_J (I \cap K_j))$. Thus by definition of $\mathcal{U}_\Phi$, we have $p \in \Phi(\bigcup_J (I \cap K_j))$, as required. The limit case is trivial.   □

THEOREM 3.5. The following are equivalent:

(1) $\Phi(S) = \Upsilon(S)$ for all $S \in \wp(S)$;
(2) $\Phi[\wp(P)]$ is a frame when considered to be a lattice ordered by inclusion.

PROOF. (1) $\implies$ (2). Let J be an indexing set and let $I, K_j \in \Phi[\wp(P)]$ for all $j \in J$. We must show that $I \cap \Phi(\bigcup_J K_j) = \Phi(\bigcup_J (I \cap K_j))$. Since we are assuming that $\Phi = \Upsilon$ and that the right side is always included in the left, it remains only to show that $I \cap \Upsilon(\bigcup_J K_j) \subseteq \Phi(\bigcup_J (I \cap K_j))$, and this is Lemma 3.4.

(2) $\implies$ (1). We note that $\Phi_0(S) = \Upsilon_0(S)$ for all $S \in \wp(P)$, by definition, and we proceed by using transfinite induction to show that $\Phi_\alpha(S) \subseteq \Upsilon_\alpha(S)$ for all cardinals α and for all $S \in \wp(P)$. As $\Upsilon(S) \subseteq \Phi(S)$ by Lemma 3.3, the result then follows.

Let $S \in \wp(P)$ and suppose $\Phi_\alpha(S) \subseteq \Upsilon_\alpha(S)$ for some cardinal α. Let $T \subseteq \Phi_\alpha(S)$ and suppose $T \in \mathcal{U}$. Let $p \leq \bigvee T$. In $\Phi[\wp(P)]$, we have $p^\downarrow \cap \bigvee_T t^\downarrow = p^\downarrow \cap (\bigvee T)^\downarrow = p^\downarrow$ by Proposition 2.9. Since $\Phi[\wp(P)]$ is a frame, $\bigvee_T(p^\downarrow \cap t^\downarrow) = p^\downarrow$, and thus $\bigvee_T(p^\downarrow \cap t^\downarrow)$ $= \Phi(\bigcup_T(p^\downarrow \cap t^\downarrow)) = \Phi(T^\downarrow \cap p^\downarrow) = p^\downarrow$. So $p = \bigvee(T^\downarrow \cap p^\downarrow)$ by Lemma 2.18(1) and $T^\downarrow \cap p^\downarrow \in \mathcal{U}_\Phi$ by Lemma 2.18(2). But $T^\downarrow \subseteq \Phi_\alpha(S) \subseteq \Upsilon_\alpha(S)$, and thus $T^\downarrow \cap p^\downarrow \subseteq \Upsilon_\alpha(S)$ and $p \in \Upsilon_{\alpha+1}(S)$, as required. The limit case is trivial and so the proof is complete. □

COROLLARY 3.6. *If $\Phi[\wp(P)]$ is not a frame, then there is $T \in \mathcal{V}$ and $p \in P$ such that*

(1) $p \leq \bigvee T$ *and*

(2) $\phi(p) \wedge \bigvee \phi[T] \neq \bigvee_T(\phi(p) \wedge \phi(t))$ *in $\Phi[\wp(P)]$.*

Here ϕ is the canonical map from P into $\Phi[\wp(P)]$ taking p to $p^\downarrow$.

PROOF. By Theorem 3.5, if $\Phi[\wp(P)]$ is not a frame, then there is $S \in \wp(P)$ with $\Phi(S) \neq \Upsilon(S)$. Let α be the smallest cardinal such that $\Phi_\alpha(S) \nsubseteq \Upsilon(S)$. This exists by Lemma 3.3 and the assumption that $\Phi(S) \neq \Upsilon(S)$. Moreover, α cannot be zero or a limit cardinal so $\alpha = \beta + 1$ for some β. Choose any $p \in \Phi_\alpha(S) \setminus \Upsilon(S)$. Then by minimality of α, we must have $p \notin \Phi_\beta(S)$, and $p \leq \bigvee T$ for some $T \in \mathcal{V}$ with $T \subseteq \Phi_\beta(S)$. In $\Phi[\wp(P)]$ we have $p^\downarrow \cap \bigvee_T t^\downarrow = p^\downarrow$. However, if $p^\downarrow = \bigvee_T(p^\downarrow \cap t^\downarrow) = \Phi(T^\downarrow \cap p^\downarrow)$, then $p = \bigvee(T^\downarrow \cap p^\downarrow)$ and $T^\downarrow \cap p^\downarrow \in \mathcal{U}_\Phi$ by Lemma 2.18. Since $T^\downarrow \cap p^\downarrow \subseteq \Phi_\beta(S)$, by choice of α it follows that $T^\downarrow \cap p^\downarrow \subseteq \Upsilon(S)$. Since $\Upsilon(S)$ is closed under joins from $\mathcal{U}_\Phi$ (by an argument similar to that in the proof of Proposition 2.8) and $p \notin \Upsilon(S)$ by choice of p, we must, therefore, have $p^\downarrow \neq \bigvee_T(p^\downarrow \cap t^\downarrow)$ to avoid a contradiction. □

Here the minimality of $\mathcal{V}$ is relevant as it gives a smaller upper bound on the possible size of a cardinal α for which α-distributivity can fail. Two immediate consequences of Corollary 3.6 are the following corollaries.

COROLLARY 3.7. *If $\Phi[\wp(P)]$ is not a frame, then $\Phi[\wp(P)]$ must fail to be σ-distributive (recall that σ is the radius of $\mathcal{V}$).*

PROOF. If $\Phi[\wp(P)]$ is not a frame, then distributivity fails for $\phi(p) \wedge \bigvee \phi[T]$ for some $T \in \mathcal{V}$, and $|T| < \sigma$ by definition. □

COROLLARY 3.8. *Given a poset P, the lattice of all downsets of P closed under existing finite joins is a frame if and only if it is distributive.*

PROOF. This lattice is produced by the closure operator which arises from the join-specification containing all nonempty finite sets with defined joins. The radius of this join-specification is ω, and the result then follows from Corollary 3.7. □

If κ is a regular cardinal and $\mathcal{V} = \mathcal{U}_\kappa = \{S \in \wp(P) \setminus \{\emptyset\} : |S| < \kappa$ and $\bigvee S$ exists$\}$, Theorem 3.5 and its corollaries are superseded by [12, Theorem 2.7]. Indeed, in this case, the mentioned theorem shows, using our notation, that $\Phi[\wp(P)]$ is a frame if and only if $\Phi(S) = \{\bigvee T : T \in \mathcal{V}$ and $T \subseteq S^\downarrow\}$ for all $S \in \wp(P)$.

Various distributivity properties for join-completions of posets (and quasiorders) are investigated in [9]. In particular, see [9, Theorem 2.1] for a summary of other properties equivalent to the join-completion corresponding to a given standard closure operator being a frame.

4. Sentences guaranteeing representability

DEFINITION 4.1 ((α, β)-representable). For cardinals α and β, a poset P is (α, β)-representable if there is an embedding $h : P \to F$, where F is a powerset algebra, such that h preserves meets of sets with cardinalities strictly less than α and joins of sets with cardinalities strictly less than β. When $\alpha = \beta$, we just write α-representable.

The next definition appears under a slightly different name as [3, Definition 2].

DEFINITION 4.2 (**MD**). A poset P is **MD** (meet distributive) if, for all $a, b, c \in P$, if $a \wedge (b \vee c)$ is defined, then $(a \wedge b) \vee (a \wedge c)$ is also defined and the two are equal.

Given a cardinal $\alpha \geq 3$, we can generalise this to the following definition.

DEFINITION 4.3 (**MD**$_\alpha$). A poset P is **MD**$_\alpha$ if, whenever $\{a\} \cup X \subseteq P$, $|X| < \alpha$ and $a \wedge \bigvee X$ is defined in P, then $\bigvee_X (a \wedge x)$ is also defined in P and the two are equal.

Note that, for semilattices, **MD**$_\alpha$ is equivalent to α-distributivity. We use our terminology to avoid confusion with the distributivity for posets defined in [12]. Given a cardinal γ, suppose that $\mathcal{V} = \mathcal{U}_\gamma$ is the set of all nonempty sets $S \subseteq P$ such that $|S| < \gamma$ and $\bigvee S$ exists in P and define Φ and Υ for this $\mathcal{V}$ as in Section 3. Note that in Section 3 we say that $\mathcal{V}$ is minimal, but minimality is not essential in the definitions of Φ and Υ or the theory developed therein. So the fact that $\mathcal{U}_\gamma$ may not be minimal is not a problem here. We have the following theorem.

THEOREM 4.4. *Let γ be any cardinal strictly greater than two. If P is **MD**$_\gamma$, then $\Phi[\wp(P)]$ is a frame.*

PROOF. By Corollary 3.6, if $\Phi[\wp(P)]$ fails to be a frame, there must be $T \in \mathcal{V}$ and $p \in P$ such that $p \leq \bigvee T$ and $\phi(p) \wedge \bigvee \phi[T] \neq \bigvee_T (\phi(p) \wedge \phi(t))$ in $\Phi[\wp(P)]$. But, by the definition of $\mathcal{V}$, we must have $|T| < \gamma$. So by Proposition 2.9 and the assumption that P is **MD**$_\gamma$, it follows that $\phi(p) \wedge \bigvee \phi[T] = \phi(p \wedge \bigvee T) = \phi(\bigvee_T (p \wedge t)) = \bigvee_T (\phi(p) \wedge \phi(t))$, so $\Phi[\wp(P)]$ must be a frame after all. $\square$

COROLLARY 4.5. *Every **MD** poset has an $(\omega, 3)$-representation.*

PROOF. Let $\mathcal{V} = \mathcal{U}_3$ and let Φ be defined as in Section 3. Since $\Phi[\wp(P)]$ is distributive, it embeds into a powerset algebra F via a map preserving finite meets and joins. Since $\phi : P \to \Phi[\wp(P)]$ preserves binary joins and arbitrary meets, we obtain an $(\omega, 3)$-representation for P by composing ϕ with the embedding of $\Phi[\wp(P)]$ into F. $\square$

By setting the value of γ appropriately, we also obtain the following result.

COROLLARY 4.6. *If $3 \leq n \leq \omega$, then every **MD**$_n$ poset has an (ω, n)-representation.*

When $\alpha > \omega$, we do *not* obtain a result corresponding to Corollary 4.6 because a frame is not necessarily (ω, α)-representable in this case. We know this because, for every $\alpha > \omega$, every nonatomic countable Boolean algebra satisfies $\mathbf{MD}_\alpha$ but is not (ω, α)-representable, as discussed in the introduction.

Note that it follows from Theorem 4.4 and [12, Theorem 2.7] that if κ is a regular cardinal, then $\mathbf{MD}_\kappa \implies \kappa$-distributive in the sense of [12] (which, for convenience, we shall dub $\mathbf{HM}_\kappa$). Also by [12, Theorem 2.7], being $\mathbf{HM}_\kappa$ is equivalent to having $\Phi(S) = \{\bigvee T : T \in \mathcal{U}_\kappa \text{ and } T \subseteq S^\downarrow\}$ for all $S \in \wp(P)$ (as mentioned following Corollary 3.8). Using this, we can show that $\mathbf{HM}_\kappa \not\Longrightarrow \mathbf{MD}_\kappa$ for all regular κ, as Example 4.7 below provides a poset that is $\mathbf{HM}_\kappa$, but fails to be $\mathbf{MD}_\kappa$, for all regular κ.

EXAMPLE 4.7. Let P be the poset in the diagram below and let κ be any regular cardinal. Given $S \in \wp(P)$, define $\Phi'(S) = \{\bigvee T : T \in \mathcal{U}_\kappa \text{ and } T \subseteq S^\downarrow\}$. Then $a \wedge (b \vee c) = c$, but $a \wedge b$ does not exist. So P fails to be $\mathbf{MD}_3$, and thus fails to be $\mathbf{MD}_\kappa$. However, the only nontrivial nonprincipal downsets of P are $\{b, c\}$ and $\{a, b, c\}$, and $\Phi(\{b, c\}) = (b \vee c)^\downarrow = \Phi'(\{b, c\})$ and $\Phi(\{a, b, c\}) = P = \Phi'(\{a, b, c\})$, and so P is $\mathbf{HM}_\kappa$ for all regular κ.
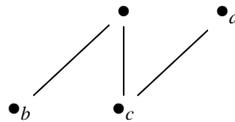
By combining the preceding discussion with [12, Example 4.2 and Theorem 4.8], we obtain the following chain of strictly one way implications.

$$\mathbf{MD}_\omega \implies \mathbf{HM}_\omega \implies \omega\text{-representable.}$$

5. Obstructions to representability

Another way to view these results is to think about obstructions to a poset being, for example, 3-representable. For the sake of this discussion, we say that P has a *triple* (a, b, c) if there are $a, b, c \in P$ such that $a \wedge (b \vee c)$ is defined. If (a, b, c) is a triple of P and $(a \wedge b) \vee (a \wedge c)$ is defined but not equal to $a \wedge (b \vee c)$, we call (a, b, c) a *split triple*, and we say that a triple (a, b, c) is an *indeterminate triple* if $(a \wedge b) \vee (a \wedge c)$ is not defined.

It is easy to see that having a split triple is a sufficient condition for a poset P to fail to be 3-representable. Moreover, existence of a split triple can be defined in first-order logic, so if having a split triple were also a necessary condition for a poset to fail to be 3-representable, it would follow that the class of 3-representable posets is finitely axiomatisable. This is not the case [6], so there must be other, less obvious obstructions to 3-representability. This is not surprising, as even in the simpler semilattice case the existence of a split triple is not necessary for failure of 3-representability. However, the semilattices that fail to be $(\omega, 3)$-representable can be characterised as those that contain either a split triple or an indeterminate triple (by [2, Theorem 2.2], phrased in the terminology of triples), which is also a

first-order property. Since the class of (m, n)-representable posets is elementary for all m and n with $2 < m, n \leq \omega$, it follows that the class of posets that fail to be 3-representable cannot be axiomatised in first-order logic at all (otherwise it would be finitely axiomatisable, in contradiction with [6]). This contrasts starkly with the intuitive finite axiomatisation of the semilattice case.

Putting this another way, let $\mathcal{S}$ be the class of posets containing a split triple, let $\mathcal{I}$ be the class of posets containing an indeterminate triple but no split triple, let $\mathcal{L}$ be the class of **MD**-posets and let $\mathcal{R}$ be the class of 3-representable posets. Then $\mathcal{S}$, $\mathcal{I}$ and $\mathcal{L}$ are all basic elementary and partition the class of all posets. Moreover, $\mathcal{L} \subset \mathcal{R}$ and $\mathcal{S} \subset \bar{\mathcal{R}}$. However, $\mathcal{I} \cap \mathcal{R}$ is elementary but not finitely axiomatisable, and $\mathcal{I} \cap \bar{\mathcal{R}}$ is not even elementary. This contrasts with the semilattice case where $\mathcal{L}$ and $\mathcal{R}$ coincide.

On the other hand, one reason why we might expect these obstructions to representability to defy simple characterisation comes from computational complexity theory. It was shown in [19] that the problem of deciding whether a finite poset has an (m, n)-representation is **NP**-complete for countable m, n with at least one of them greater than three. So if there were a simple characterisation of the finite posets that fail to be (m, n)-representable, such as exists in the semilattice case, we could potentially use this to prove that this decision problem is also in **coNP**, which would imply the unexpected coincidence **NP** = **coNP**. More precisely, by Fagin's theorem, **NP** = **coNP** if and only if the finite posets that fail to be (m, n)-representable (for suitable m and n) have an existential second-order characterisation.

References

[1]   A. Abian, 'Boolean rings with isomorphisms preserving suprema and infima', *J. Lond. Math. Soc. (2)* **3** (1971), 618–620.

[2]   R. Balbes, 'A representation theory for prime and implicative semilattices', *Trans. Amer. Math. Soc.* **136** (1969), 261–267.

[3]   Y. Cheng and P. Kemp, 'Representation of posets', *Zeitschr. f. math. Logik und Grundlagen d. Math* **38** (1992), 269–276.

[4]   W. Cornish and R. Hickman, 'Weakly distributive semilattices', *Acta Math. Acad. Sci. Hungar.* **32** (1978), 5–16.

[5]   H. P. Doctor, 'Extensions of a partially ordered set', PhD Thesis, McMaster University (Canada), 1967.

[6]   R. Egrot, 'No finite axiomatizations for posets embeddable into distributive lattices', Preprint, 2016, arXiv:1610.00858.

[7]   R. Egrot, 'Representable posets', *J. Appl. Log.* **16** (2016), 60–71.

[8]   M. Erné, 'Closure', in: *Beyond Topology*, Contemporary Mathematics, 486 (Amer. Math. Soc., Providence, RI, 2009), 163–238.

[9]   M. Erné and G. Wilke, 'Standard completions for quasiordered sets', *Semigroup Forum* **27**(1–4) (1983), 351–376.

[10] G. Grätzer and E. T. Schmidt, 'On congruence lattices of lattices', *Acta Math. Acad. Sci. Hungar.* **13** (1962), 179–185.

[11] R. Hickman, 'Distributivity in semilattices', *Acta Math. Acad. Sci. Hungar.* **32** (1978), 35–45.

[12] R. C. Hickman and G. P. Monro, 'Distributive partially ordered sets', *Fund. Math.* **120** (1984), 151–166.

[13] K. Kearnes, 'The class of prime semilattices is not finitely axiomatizable', *Semigroup Forum* **55** (1997), 133–134.

[14] P. Kemp, 'Representation of partially ordered sets', *Algebra Universalis* **30** (1993), 348–351.
[15] B. Schein, 'On the definition of distributive semilattices', *Algebra Universalis* **2** (1972), 1–2.
[16] J. Schmidt, 'Universal and internal properties of some completions of k-join-semilattices and k-join-distributive partially ordered sets', *J. reine angew. Math.* **255** (1972), 8–22.
[17] J. Schmidt, 'Each join-completion of a partially ordered set is the solution of a universal problem', *J. Aust. Math. Soc.* **17** (1974), 406–413.
[18] J. Schmidt, 'Universal and internal properties of some extensions of partially ordered sets', *J. reine angew. Math.* **253** (1972), 28–42.
[19] C. Van Alten, 'Embedding ordered sets into distributive lattices', *Order* **33** (2016), 419–427.
[20] J. C. Varlet, 'On separation properties in semilattices', *Semigroup Forum* **10** (1975), 220–228.

ROB EGROT, Faculty of Information and Communication Technology,
Mahidol University, 999 Phuttamonthon 4 Road, Salaya,
Nakhon Pathom 73170, Thailand
e-mail: robert.egr@mahidol.ac.th