Modern faces of filtration

Johan van Benthem and Nick Bezhanishvili

This paper is dedicated to Kit Fine, a prominent pioneer in mathematical modal logic. Several of his classical themes shine through in the course of what follows.

Abstract The filtration method for proving decidability in a focused minimal manner is a highlight of modal logic, widely used, but also posing a bit of a challenge as to its scope and what makes it tick. In this paper, we bring together a number of modern perspectives on filtration, including model-theoretic and proof-theoretic ones. We also include a few more unusual recent connections with dynamic logics of model change and logics of questions and issues. Finally, we analyze where the filtration method fails in full first-order logic, and what it still has to say there.

1 Introduction

The classical filtration method is a typical modal tool, widely used, but also posing a bit of a challenge to understand what makes it tick. It has been used to prove decidability of a wide variety of modal logics, couched in a variety of formal languages, and its field of application is stil expanding. We will provide some general perspectives on the filtration method, drawing mainly on recent literature, and adding a few new observations of our own. We hope that this presentation is useful, since no such story seems to exist so far. Even so, this is not a self-contained introduction. The reader is assumed to have some basic knowledge of modal logic and filtration.

Johan van Benthem
Institute for Logic, Language and Computation, University of Amsterdam, Postbus 94242, 1090 GE Amsterdam, The Netherlands. e-mail: `j.vanbenthem@uva.nl` & Amsterdam-Beijing Joint Research Center in Logic, Tsinghua University.

Nick Bezhanishvili
Institute for Logic, Language and Computation, University of Amsterdam, Postbus 94242, 1090 GE Amsterdam, The Netherlands. e-mail: `N.Bezhanishvili@uva.nl`

In Section 2, the original basic results about filtration are stated, followed by impressions of the state of the art in filtration for a range of modal logics, and a range of modal languages. The section concludes with a discussion of what a filtration really is, a question to which new answers keep emerging. In Section 3, more constructive versions of filtration are presented that establish completeness at the same time, by now perhaps the method favored by working modal logicians. This style relates to the use of normal forms and issues of computational complexity for modal logics, that will be discussed as well.

After this more survey-style first part, we turn to new topics and results. Section 4 presents a recent development: a dynamic modal logic that axiomatizes the basic properties of filtration, thus drawing the meta-theory of filtration into modal logic. Behind this logic is the view that filtration is one member of a wider family of natural notions of model change, and in Section 5, the interaction is analyzed between filtration and other modal logics of model change, in particular, dynamic-epistemic logic of information update. A further general aspect of filtration is discussed in Section 6: its relation to choosing a set of issues that focus inquiry. Connections are found with logics of questions and issues, and as a result, a new open problem is identified about the complete logic of filtration with varying vocabularies. In Section 7, the boundaries of filtration are discussed, since the method clearly does not work for first-order logic, which is undecidable. However, filtration still works partially in the first-order realm, and connections are brought to light with generalized semantics for decidable versions of first-order logic.

In all, then, the reach of filtration is wider than might be thought. Section 8 wraps up and suggests further directions arising from looking at filtration in its full generality. Many of these directions are technical, but there is also an interesting conceptual issue. Many philosophical models of partiality restrict an atomic vocabulary, but then consider any formula in that vocabulary. In contrast with this, filtration is 'super-partial': it also restricts the shapes of assertions that are considered. We believe that this makes it of wider interest than just a technical tool for decidability.

2 Basic theory of filtration

2.1 The initial paradigm

We start by recalling the model-theoretic approach to filtrations (see, e.g., [25, Sec. 2.3] or [29, Sec. 5.3]). Let $\mathfrak{M} = (W, R, V)$ be a relational model and let Σ be a set of formulas closed under subformulas. For our purposes, Σ will always be assumed to be finite. Define an equivalence relation $\sim_\Sigma$ on the set of worlds W by

$$x \sim_\Sigma y \text{ iff } (\forall \varphi \in \Sigma)(\mathfrak{M}, x \models \varphi \Leftrightarrow \mathfrak{M}, y \models \varphi).$$

We introduce a new model structure $\mathfrak{M}^f = (W^f, R^f, V^f)$. Let $W^f = W/\!\sim_\Sigma$ and let $V^f(p) = \{[x] : x \in V(p)\}$, where $[x]$ is the equivalence class of x with respect to $\sim_\Sigma$.[1]

Definition 2.1 *For a binary relation R^f on W^f, the model $\mathfrak{M}^f = (W^f, R^f, V^f)$ is a* filtration of $\mathfrak{M}$ through Σ *if the following two conditions are satisfied:*[2]

(F1) $xRy \Rightarrow [x]R^f[y]$.

(F2) $[x]R^f[y] \Rightarrow (\forall \Diamond\varphi \in \Sigma)(\mathfrak{M}, y \models \varphi \Rightarrow \mathfrak{M}, x \models \Diamond\varphi)$.

Evidently, if Σ is finite, then the new set of worlds W^f is finite. In fact, if Σ consists of n elements, then W^f consists of no more than 2^n elements.

There are several ways of meeting the filtration conditions. In particular, the 'smallest filtration' of $\mathfrak{M} = (W, R, V)$ through Σ is $\mathfrak{M}^s = (W^f, R^s, V^f)$ and the 'largest filtration' is $\mathfrak{M}^l = (W^f, R^l, V^f)$, where

$$[x]R^s[y] \text{ iff } (\exists x', y' \in W)(x \sim_\Sigma x' \ \&\ y \sim_\Sigma y' \ \&\ x'Ry').$$
$$[x]R^l[y] \text{ iff } (\forall \Diamond\varphi \in \Sigma)(\mathfrak{M}, y \models \varphi \Rightarrow \mathfrak{M}, x \models \Diamond\varphi).$$

If $\mathfrak{M}^f = (W^f, R^f, V^f)$ is a filtration of $\mathfrak{M} = (W, R, V)$, then indeed $R^s \subseteq R^f \subseteq R^l$.

Theorem 2.2 (Filtration Theorem) *Let $\mathfrak{M}^f = (W^f, R^f, V^f)$ be a filtration of the model $\mathfrak{M} = (W, R, V)$. Then for every formula $\varphi \in \Sigma$ and $w \in W$, we have*

$$\mathfrak{M}, w \models \varphi \textit{ iff } \mathfrak{M}^f, [w] \models \varphi.$$

This fundamental fact follows by a simple induction on the formula φ. It implies that Basic Modal Logic has the *Finite Model Property* (often abbreviated FMP in what follows): if a formula is satisfiable at all, it is satisfiable in some finite model.

Corollary 2.3 (Decidability) *Satisfiability in the Basic Modal Logic is decidable.*

A brute force decision method relying on Theorem 2.2. starts from a given formula φ, and checks if φ is true at any point in any model whose domain has at most $2^{|Sub(\varphi)|}$ points. In Section 3, we return to the actual complexity of satisfiability.

This argument depends on the fact that the relevant finite models form a finite set, since the filtration method as stated above gave an upper bound on the size of the models. This is sometimes called the 'Effective Finite Model Property'. Without it, a modal logic can have the FMP and still be undecidable, [64]. We will discuss further potential subtleties with decidability below.

Filtration in this sense is purely model-theoretic, working on models.

---

[1] The valuation makes most sense if p occurs in Σ, but we could also interpret the whole language by making the proposition letters p outside of Σ uniformly true or false.

[2] Strictly speaking, we should indicate the set Σ used in creating the filtrated model, but an explicit notation $\mathfrak{M}^f_\Sigma$ will not be needed, since the relevant Σ will usually be clear from the context.

Variants of the method are used when more structure needs to be preserved. For instance, a filtration of a transitive model may not be transitive. Therefore, we introduce two further notions of filtration (see, e.g., [29, Sec. 5.3] or [25, Sec. 2.3]).

The *transitive or Lemmon filtration* of $\mathfrak{M} = (W, R, V)$ through Σ is given by $[x]R^t[y]$ iff $(\forall \Diamond\varphi \in \Sigma)(\mathfrak{M}, y \models \varphi \vee \Diamond\varphi \Rightarrow \mathfrak{M}, x \models \Diamond\varphi)$. The smallest transitive filtration relation R^{st} is obtained by the transitive closure of the smallest filtration. It is well known that if $\mathfrak{M} = (W, R, V)$ has R transitive, then $\mathfrak{M}^t = (W^f, R^t, V^f)$ and $\mathfrak{M}^{st} = (W^f, R^{st}, V^f)$ are filtrations of $\mathfrak{M}$ whose relations R^t and R^{st} are transitive.

The filtration method as presented here follows the model-theoretic line of the pioneering publications [56, 62]. However, there are also early algebraic perspectives on filtration, [58, 59], and we discuss some modern versions in Section 3.5 below.

Remark 2.4 There are in fact two common methods of proving the finite model property for modal logics: filtration or 'standard filtration' and 'selective filtration'. The method of selective filtration was developed by Dov Gabbay in [39, 42, 44] and, in a version for transitive modal logics, by Kit Fine [37]. The latter version also led to new axiomatization methods (see the next section for more details) and to the important notions of 'subframe logics' and 'cofinal subframe logics' (see [29, Ch. 9] for a comprehensive overview).

The filtration technique has proved remarkably extendable, in two directions:

2.2 Stronger modal axioms

Staying within the standard modal language, extended filtration methods guarantee more frame properties for logics with additional axioms. Given a normal modal logic L (an extension of the Basic Modal Logic K), the filtration theorem gives us that, if a filtration of an L-model results in a model based on a frame of L (such logics are said to *admit filtration*), then L has the FMP. This yields the FMP for many well-known modal logics e.g., K, KT, KB, $\mathsf{K4}$, $\mathsf{S4}$, $\mathsf{S4.2}$ and many more. We refer to [25] and [29, Sec. 5.3] for the details. An interesting example is the modal logic $\mathsf{S4.3}$. Not only does it enjoys the FMP, which can easily be shown via filtration, but Bull [27] algebraically, and Fine [35] model-theoretically, proved that every normal extension of $\mathsf{S4.3}$ also has the FMP. In fact, Fine's proof provides an interesting synergy of filtration and selection, [35], see also [25, Sec. 4.9]. An alternative proof is given in [67], see also [29, Ch 11], where it is shown that every normal extension of $\mathsf{S4.3}$ is a cofinal subframe logic.

The above definition of a logic admitting filtration involves many parameters: models, frames, filtrations. In an attempt to obtain a uniform theory of filtrations, *stable logics* were introduced in [18] and investigated further in [19] and [51] as modal logics whose class of rooted frames is closed under order-preserving images. [19] and [51] also study $\mathsf{K4}$ and $\mathsf{S4}$-*stable logics* (more generally *L-stable logics*). Every stable, $\mathsf{K4}$-stable or $\mathsf{S4}$-stable logic admits filtration in the above sense and hence enjoys the FMP. In a way stable logics are those logics that admit all filtrations (all transitive filtrations in case of $\mathsf{K4}$-stable or $\mathsf{S4}$-stable logics).

However, filtration and selective filtration are not only techniques for proving the FMP—they can also be turned into axiomatization methods. For transitive modal logics this approach via selective filtration was pioneered by Kit Fine [35, 37]. With each finite transitive rooted frame F, Fine associated its subframe formula which is refuted on a frame G iff there is a subframe of G bounded morphically mapped onto F. Then every subframe logic above $\mathsf{K4}$, i.e., a logic whose frame class is closed under taking subframes, is axiomatized by such formulas [37]. Canonical formulas were introduced in [66] (see also [29, Ch. 9]) as generalizations of subframe formulas and it was shown that every logic above $\mathsf{K4}$ is axiomatized by canonical formulas. The proof of this result relies essentially on a suitably general version of selective filtration.

A parallel approach of axiomatizing modal logics via filtration was undertaken recently. For each finite transitive rooted frame one can define its *stable formula* and show that every $\mathsf{K4}$-stable logic is axiomatizable by such formulas [18, 19]. Moreover, for each finite rooted frame and a fixed set of subsets of this frame, one defines its *stable canonical rule*. If this frame is transitive, then this rule is equivalent to a formula, called a *stable canonical formula*.

To illustrate this dense survey, we display one typical modern result. For a proof of the following theorem, using filtrations in an essential way, see [18].

Theorem 2.5 *Every normal modal logic is axiomatizable by stable canonical rules. Moreover, every normal modal logic above* $\mathsf{K4}$ *is axiomatizable by stable canonical formulas and every* $\mathsf{K4}$*-stable logic is axiomatizable by stable formulas.*

This result yields immediately that there exist continuum many stable logics and continuum many $\mathsf{K4}$ and $\mathsf{S4}$-stable logics [19, 51]. All these logics admit filtration. Since there are only countably many decidable logics, uncountably many stable logics must be undecidable. Thus, there are uncountably many undecidable modal logics that admit filtration and hence have the FMP with an effective bound. Crucially, none of these logics can be finitely axiomatizable – the earlier simple argument deriving decidability from effectively bounded filtration breaks down as checking the truth of infinitely many axioms, even in a given finite space of potential countermodels, is not an effective decision method. However, each logic axiomatized by finitely many stable formulas will admit filtration and hence be decidable.

2.3 Richer modal languages

The second direction in which filtration methods have been generalized concerns extensions to richer languages. We start with a simple case, where we just add a universal modality to the basic modal language.

Fact 2.6 *The filtration theorem also holds with the universal modality added.*

The proof is similar to the proof of Theorem 2.2, except that a straightforward additional inductive step for universal formulas needs to be verified.

In general, filtration sets may have to be chosen with great care and imagination, when a modal language gets extended. We briefly discuss an influential case: Propositional Dynamic Logic PDL, a running example in several sections to follow.

Definition 2.7 *The* dynamic language PDL *is defined by the grammar:*

$$\pi := a \mid ?\varphi \mid \pi;\pi \mid \pi\cup\pi \mid \pi^* \quad \textit{and}$$
$$\varphi := p \mid \neg\varphi \mid \varphi\wedge\varphi \mid \langle\pi\rangle\varphi,$$

Here, a is an element of the set of basic programs Π_0. Note the mutually recursive set-up of this syntax with formulas and program expressions on a par.

Definition 2.8 *A* PDL*-model is a standard model for a poly-modal language, i.e., a tuple $\mathfrak{M} = (W, \{R_a\}_{a\in\Pi_0}, V)$, where each R_a is a binary relation.*

To define the semantics, two things must happen in tandem: giving the denotations of program expressions and of formulas. For program expressions we set:

$$R_{\pi_1\cup\pi_2} := R_{\pi_1} \cup R_{\pi_1}$$
$$R_{\pi_1;\pi_2} := R_{\pi_1} \circ R_{\pi_2}$$
$$R_{\pi^*} := (R_\pi)^*$$
$$R_{?\varphi} := \{(s,s) \mid \mathfrak{M}, s \models \varphi\}$$

Note that this depends on the truth definition of formulas in the clause for tests. Simultaneously, therefore, we define truth of a PDL-formula φ in a model $\mathfrak{M}$, which can be done in the standard way, noting that the clause for the general program modality depends on our semantics for program expressions:

$$\mathfrak{M}, w \models \langle\pi\rangle\varphi \quad \text{iff} \quad \text{there is } v \text{ such that } wR_\pi v \text{ with } \mathfrak{M}, v \models \varphi.$$

Definition 2.9 *Let Σ be a set of* PDL*-formulas. Then Σ is* Fisher-Ladner closed *if it is closed under subformulas and under taking single negations,*[3] *and which satisfies the following additional closure conditions:*

1. *If $\langle\pi_1;\pi_2\rangle\varphi \in \Sigma$, then $\langle\pi_1\rangle\langle\pi_2\rangle\varphi \in \Sigma$.*
2. *If $\langle\pi_1\cup\pi_2\rangle\varphi \in \Sigma$, then $\langle\pi_1\rangle\varphi \vee \langle\pi_2\rangle\varphi \in \Sigma$.*
3. *If $\langle\pi^*\rangle\varphi \in \Sigma$, then $\langle\pi\rangle\langle\pi^*\rangle\varphi \in \Sigma$.*

The Fisher-Ladner closure of Σ *is the smallest set of formulas containing Σ that is Fisher-Ladner closed and is denoted by $FL(\Sigma)$.*

It is a well-known fact that the Fisher-Ladner closure of a finite set Σ is still finite.

Fact 2.10 *If a* PDL*-formula φ is satisfiable in a* PDL*-model, then it is also satisfiable in a finite* PDL*-model.*

[3] Double negations $\neg\neg\psi$ are identified here with the original formulas ψ.

Proof. Let φ be a formula satisfied in a PDL-model $\mathfrak{M} = (W, \{R_a\}_{a \in \Pi_0}, V)$. We take the smallest filtration $\mathfrak{M}^f = (W^f, \{R^s_a\}_{a \in \Pi_0}, V')$ of $\mathfrak{M}$ via $FL(\Sigma)$, for $\Sigma = \{\varphi\}$. By the above observation, W^f is finite. Recall that, in this model, we have $[w]R^s_a[v]$ iff there are $w' \in [w]$ and $v' \in [v]$ such that $w' R_a v'$.

To prove the filtration theorem, we show two things in a simultaneous induction, for all programs π and formulas φ:

- For any two $w' \in [w]$ and $v' \in [v]$ with $w' R_\pi v'$, we also have $[w]R^s_\pi[v]$
- $\mathfrak{M}, w \models \varphi$ iff $\mathfrak{M}^f, [w] \models \varphi$.

The first fact is a straightforward induction on program operations, where the case of test programs appeals to the second clause. As for the second clause, we display only the crucial case when ψ has the form $\langle\pi\rangle\psi$.

Claim 2.11 *For every $w \in W$ and $\langle\pi\rangle\psi \in FL(\Sigma)$, we have*

$$\mathfrak{M}, w \models \langle\pi\rangle\psi \quad \textit{iff} \quad \mathfrak{M}^f, [w] \models \langle\pi\rangle\psi.$$

From left to right, the proof is immediate by an appeal to the first clause about preservation of program relations under filtration and the inductive hypothesis.

From right to left, we show by induction on the complexity of programs π that the assertion holds for each π with respect to all formulas φ in the Fisher-Ladner set. The case when $\pi \in \Pi_0$ is as for the standard filtration theorem. Thus, we only need to consider the cases (1) $\pi = \pi_1 \cup \pi_2$, (2) $\pi = \pi_1 ; \pi_2$, and (3) $\pi = \rho^*$.

Cases (1), (2) involve straightforward appeals to the inductive hypothesis. Now, let $\mathfrak{M}^f, [w] \models \langle\rho^*\rangle\psi$. Then there are $w_1, \ldots, w_n$ with $[w] = [w_0]R^s_\rho[w_1]R^s_\rho \ldots R^s_\rho[w_n]$, and $\mathfrak{M}^f, [w_n] \models \psi$. By the inductive hypothesis, $\mathfrak{M}, w_n \models \psi$. Then also $\mathfrak{M}, w_n \models \langle\rho^*\rangle\psi$ – which takes care of sequences of length 1. Longer finite sequences can be treated as finite relational compositions, noting that we do not need to talk about truth of ever larger iterated formulas $\langle\pi\rangle \ldots \langle\pi\rangle\psi$ (which may not be in the Fisher-Ladner set) along the sequence, but can stick with formulas $\langle\rho^*\rangle\psi$ – using the fact, in each step, that $\langle\rho\rangle\langle\rho^*\rangle\psi$ is in the Fisher-Ladner set and implies $\langle\rho^*\rangle\psi$.[4] □

Many extended modal languages interact with filtration in interesting ways. For instance, a 'transfer result' in [24] says that, if a modal logic L admits filtration, then various desirable meta-properties of L also hold for a suitable defined hybrid companion to L adding expressive devices from hybrid logic.

[4] It should be pointed out that in the above proof, while for an atomic program a it is the case that $[w]R^s_a[v]$ implies that there are $w' \in [w]$ and $v' \in [v]$ such that $w' R_a v'$, this may no longer be true for an arbitrary program π. This can be shown with an easy example involving $\pi = a^*$. Thus, in this sense, the filtrated relations correspond to the smallest filtration only for basic programs.

2.4 Coda

But what *is* filtration? What we really require from the filtrated model is just that it satisfies the filtration theorem. This can be considered as a definition, and going further, we can even replace condition (F1) with the condition:

(F1′) $\quad \mathfrak{M}, x \models \Diamond\varphi \Rightarrow \exists y\ ([x]R^f[y]\ \&\ \mathfrak{M}, y \models \varphi).$

This condition is again equivalent to the following symmetric version of (F2):

(F1″) $\quad xRy \Rightarrow (\forall \Diamond\varphi \in \Sigma)(\mathfrak{M}^f, [y] \models \varphi \Rightarrow \mathfrak{M}^f, [x] \models \Diamond\varphi).$

Thus, by considering (F1″) and (F2), the definition of filtration can be made fully symmetric. We call this new notion *weak filtration*. It is easy to check that any of (F1′) and (F1″), together with (F2), guarantees the filtration theorem to hold. In fact, any of (F1′) and (F1″) is equivalent to the left to right direction of the $\Diamond$-clause of the filtration theorem. On the other hand, condition (F2) is equivalent to the right to left direction of the $\Diamond$-clause of the filtration theorem.[5]

One can naturally ask whether there are any logical systems which admit weak filtration but not the standard one. The answer to this question is positive. The well-known systems of modal logic GL, $\mathsf{K4.Grz}$ and $\mathsf{S4.Grz}$ all admit weak filtration in the preceding sense, but they do not admit standard filtration [26].

Weak filtration also appears to be a natural concept from a co-algebraic point of view. In co-algebraic modal logic, the theory map restricted to a subformula-closed set of formulas can be seen to satisfy a certain universal property. For the type of co-algebras that correspond to Kripke models, the filtration condition (F1′) is then precisely what is needed for a diagrammatic proof of the Truth Lemma.[6]

But there are also drawbacks to this new notion of filtration. There need not be a weakest filtration – and more importantly, we lose the appealing quotient intuition.

2.5 Conclusion

Filtration is a wide-ranging method, and its range still continues to expand.[7]

---

[5] In this format, filtration looks like a form of *bisimulation* restricted to back-and-forth behavior w.r.t. only a finite set of relevant formulas – a perspective that we will not explore further here.

[6] This point is due to personal communication with Clemens Kupke and Jurriaan Rot.

[7] Filtration also works on other similarity types than the relational models for modal logic discussed here. For instance, filtration on neighborhood semantics can be found in [30, 48, 61, 63].

3 Constructive versions of filtration

3.1 Semantic or proof-theoretic

Filtration can be a purely semantic method working on models. But it also provides results on proof-theoretic canonical Henkin models, yielding FMP and decidability plus completeness. In this section, we consider versions of filtration over syntactic Henkin models that have come into wide use. While a new term might be preferable in this setting of proof and consistency, we follow practice in the field and just speak about ‘filtration’, as the context will always disambiguate which sense is meant.

To demonstrate the above ideas, we will first give a warm-up result. We say that a logic L *enjoys the Finite Model Property* (the FMP) if there is a class C of finite frames validating the formulas in L such that, for each formula φ, we have

$$L \vdash \varphi \text{ iff } \mathsf{C} \models \varphi.$$

It is easy to see that the above is equivalent to the fact that every L-consistent formula φ is satisfied in a model based on an L-frame in C.

Fact 3.1 *The Basic Modal Logic* K *is complete and has the FMP.*

Proof. The FMP part of this result has already been proved purely semantically in Section 2.1. We sketch the main argument in the current setting. Let φ be a K-consistent formula. Then the canonical model $\mathfrak{M}^C_{\mathsf{K}}$ satisfies φ. Obviously, as all relational frames are K-frames, $\mathfrak{M}^C_{\mathsf{K}}$ is based on a K-frame. We now take the filtration of $\mathfrak{M}^C_{\mathsf{K}}$ via the set $\Sigma = Sub(\varphi)$ of subformulas of φ. By Theorem 2.2, the filtrated model satisfies φ and is clearly based on a K-frame. This finishes the proof. $\square$

Remark 3.2 Our presentation followed what most modal logicians have in mind: they start from the complete infinite Henkin model that they know and love, and filtrate it down to a finite model. However, there is an alternative which is also in wide use. One restricts the ‘language’ of the whole Henkin construction to the finite filtration set, and thus immediately obtains a finite Henkin model – noting that the usual completeness proof goes through in this rectricted setting. This may not yet be the model one wants, so further surgery may be needed, as we shall see below.

In the current style, filtration connects with the analysis of proof principles for modal logics that support its success. The above proof generalizes to every modal logic L that is canonical (i.e., the canonical model of L is based on a frame for L) and admits filtration. However, the path towards the FMP via filtration may work even for logics that are not canonical. Sometimes, the canonical frame of a logic is not in a ‘right’ class, but its filtrations are, yielding FMP and decidability after all. We will demonstrate this for the case of Propositional Dynamic Logic PDL, discussed earlier in a purely semantic setting. Other non-canonical systems for which this method works are the well-known modal logics GL, $\mathsf{K4.Grz}$ and $\mathsf{S4.Grz}$, but in these cases one needs to work with the notion of weak filtration discussed in Section 2.4.

Recall that the system PDL is axiomatized by the following principles:

(CPC)	Axioms and rules of classical propositional logic
(Ax-K_π)	$[\pi](\varphi \to \psi) \to ([\pi]\varphi \to [\pi]\psi)$
(Ax-$(?\varphi)$)	$[?\varphi]\psi \leftrightarrow (\varphi \to \psi)$
(Ax-$(\pi;\pi')$)	$[\pi;\pi']\varphi \leftrightarrow [\pi][\pi']\varphi$
(Ax-$(\pi \cup \pi')$)	$[\pi \cup \pi']\varphi \leftrightarrow ([\pi]\varphi \wedge [\pi']\varphi)$
(FIX)	$[\pi^*]\varphi \leftrightarrow (\varphi \wedge [\pi][\pi^*]\varphi)$
(IND)	$[\pi^*](\varphi \to [\pi^*]\varphi) \to (\varphi \to [\pi^*]\varphi)$
(Nec$_{[\pi]}$)	if φ is provable, then $[\pi]\varphi$ is provable

Table 1: Axioms and rules of PDL

Theorem 3.3 *Propositional Dynamic Logic* PDL *is complete and has the FMP.*

Proof. We show that, if a formula φ is PDL-consistent, then it is satisfiable.

For convenience, we follow the second view outlined above, and just consider the finite Henkin model $\mathfrak{M}^C = (W^C, \{R_a^C\}_{a \in \Pi_0}, V^C)$ consisting of all maximally consistent sets where the language of available formulas is just the filtration set.[8] Note that we can view the conjunctions of all formulas in these sets Σ as single formulas $\hat{\Sigma}$.

Next, in defining the relations $\Sigma R_a^C \Delta$ for atomic programs, we still refer to provability and consistency in the complete language beyond the filtration set:

$$\Sigma R_a^C \Delta \text{ iff } \hat{\Sigma} \wedge \langle a \rangle \hat{\Delta} \text{ is consistent in the logic PDL.}$$

More generally, we define relations $\Sigma S_\pi \Delta$ as consistency of $\hat{\Sigma} \wedge \langle \pi \rangle \hat{\Delta}$. However, we also define a second family of relations, starting from the atomic relations R_a^C and then lifting these to R_π for all programs π via the usual semantics for PDL. We also use the notation $\mathfrak{M}^C$ for the standard model arising in this second way.

Claim 3.4 *For all programs π and formulas ψ occurring in $FL(\Sigma)$, and all maximally consistent sets Δ in the finite Henkin model $\mathfrak{M}^C$, we have*

1. $S_{\pi*} \subseteq (S_\pi)^*$
2. $S_\pi \subseteq R_\pi$

Claim 3.5 *For all programs π and formulas ψ occurring in $FL(\Sigma)$, and all maximally consistent sets Δ in the finite Henkin model $\mathfrak{M}^C$, we have*

$$\psi \in \Delta \;\textit{ iff }\; \mathfrak{M}^C, \Delta \models \psi.$$

These three claims are proved by a simultaneous induction on formulas and programs, where the second claim is the one that appeals essentially to the Induction Axiom IND. In following the steps, one will use each direction of the PDL-axioms once, as well as each clause in the definition of $FL(\Sigma)$. For details, we refer to [25].

[8] This sparse approach bypasses non-standard infinite consistent sets of formulas such as $\{\langle a^* \rangle \varphi, \neg p, [a]\neg p, [a]^2\neg p, \ldots,\}$ that do not admit of an interpretation in standard models.

From the claims we obtain the desired finite model that satisfies φ. Thus, on top of its completeness, PDL has the FMP. $\square$

Remark 3.6 The following observations throw some further light on the preceding. Let W be the set of all maximal PDL-consistent sets. Then it is immediate that $W^C = \{\Gamma \cap FL(\Sigma) : \Gamma \in W\}$. For each program π, we define the canonical relation $\mathfrak{R}^C_\pi$ on W by setting $\Gamma \mathfrak{R}^C_\pi \Delta$ if, for each formula φ, $\varphi \in \Delta$ implies $\langle\pi\rangle\varphi \in \Gamma$. Next we show that for each $\Gamma, \Delta \in W$ with $\Gamma \mathfrak{R}^C_\pi \Delta$, we have $AS_\pi B$, where $A = \Gamma \cap FL(\Sigma)$ and $B = \Delta \cap FL(\Sigma)$. To see this, note that $\widehat{B} \in \Delta$. So $\langle\pi\rangle\widehat{B} \in \Gamma$. Hence, $\widehat{A} \wedge \langle\pi\rangle\widehat{B} \in \Gamma$, which implies, by consistency of Γ that $\widehat{A} \wedge \langle\pi\rangle\widehat{B}$ is consistent. So we have $AS_\pi B$. But then, the finite canonical model can be seen as the smallest filtration of the full canonical model of PDL; just identify the equivalence class $[\Gamma]$ with $\Gamma \cap FL(\Sigma)$.[9]

Question Can we use the filtration proof to automatically find the axioms for PDL, just as the Henkin proof suggests the axioms for the minimal modal logic?

We continue with a discussion of some further aspects of this setting.

3.2 Normal forms

The syntactic aspect of filtration may also be connected to the use of normal forms in modal logic, discovered by Fine [36] and revisited many times by different authors, e.g., Bellissima [5], Ghilardi [45] and Moss [60].

For example, in the basic modal logic K normal forms of formulas with variables in $P_n = \{p_0, \ldots, p_n\}$ of degree 0 are formulas of the form

$$\bigwedge_{p_i \in T} p_i \wedge \bigwedge_{p_i \in P_n \setminus T} \neg p_i$$

for $T \subseteq P_n$. Normal forms of degree $k+1$ are formulas in the following form, for $T \subseteq P_n$ and S a set of normal forms of degree k:

$$\bigwedge_{p_i \in T} p_i \wedge \bigwedge_{p_i \in P_n \setminus T} \neg p_i \wedge \bigwedge_{\psi \in S} \Diamond\psi \wedge \Box \bigvee S.$$

Remark 3.7 These normal forms are also highly reminiscent of model description formulas for given pointed models up to n-bisimulation, [60]. In that setting, one starts with enumerating the true literals at the point, and then, inductively, at each next level $\alpha+1$, describes the types of successors of the point that can be formulated at description level α.

[9] The above also shows that the finite canonical model W^C with the relations S_π is the least filtration of the (big) canonical model. However, as $S_\pi \subseteq R_\pi$, the 'real' finite canonical model is also a filtration of the big canonical model. Condition (F1) follows from the fact that $S_\pi \subseteq R_\pi$, and condition (F2) follows from the truth lemma for canonical models and the above Claim 3.5.

One can think of normal forms as follows. Finite canonical models for a given modal logic L are essentially filtrations of the canonical model for L via the finite set of formulas of fixed finite modal depth k. Moreover, 'model description formulas', a finite-depth variant of modal Scott sentences describing given pointed models up to n-bisimulation, then describe worlds in these finite canonical models. Since every formula has finite modal depth, it corresponds to a subset of the finite canonical model, and its normal form is a finite disjunction of model description formulas.

In algebraic terms, finite canonical models correspond to Lindenbaum-Tarski algebras of formulas of finite modal depth. Note that these algebras are not modal algebras in the standard sense. Applying a modal operator, say $\Box$, to a formula of modal depth less than k gives us a formula of modal depth less than $k+1$. Thus, for a modal logic L if we let B_k be the algebra of formulas of modal depth less than or equal to k, modulo L-equivalence, then we can view the modal operator $\Box$ as a meet-preserving map from B_k into B_{k+1}. Then the limit of the sequence of these partial algebras is the Lindenbaum-Tarski algebra of a given logic L. Model description formulas can be seen as formulas describing atoms of such B_k's and as each B_k is finite, every formula of modal depth k will be a finite disjunction of these formulas, providing Fine normal forms. We refer to [22, 23, 45, 60] for more details on this algebraic and coalgebraic analysis of finite canonical models and normal forms.

The connection of normal forms and finite canonical models with the so-called universal models of modal and intermediate logics [21, 29] remains to be clarified. It is known that, for transitive modal logics, the model description formulas define exactly the singletons of universal models (atoms of the corresponding free algebras) leading to an alternative description of Jankov-Fine formulas [5, 29]. A similar characterization in the intuitionistic setting via join-irreducible elements gives rise to Jankov-de Jongh formulas for intuitionistic logic [21].

Remark 3.8 The Henkin model is one unique counter-model for all non-valid formulas. In fact, for modal formulas of depth k, it suffices to take the finite k-cut-off Henkin model. The semantic FMP seemed to go model by model. However, we only need finitely many types of finite counter-model: their disjoint union, too, is a universal counter-model, and we can even contract equivalent types via bisimulation inside that model. It is easy to show, using the completeness theorem, that this semantically generated model is in fact isomorphic to the finite Henkin model.

More can be said on the topic of normal forms and model description formulas, also for our running example PDL, but we ignore this here, cf. [7, Section 5].

3.3 Complexity

The proof-theoretic road to the FMP also suggests a comparison with what is known about the computational complexity of satisfiablility and validity for modal logics. For instance, for the minimal modal logic, the satisfiability problem is known to be *Pspace*-complete. But the filtration method as explained in Section 2.1 gives an exponential size for counter-models, and hence it overshoots widely.

The reason for the mismatch is that we do not need to display counter-models explicitly in order to test for their existence. A finer analysis is via *semantic tableaux* which faithfully test the existence of satisfying models for a given formula branch by branch, [38], and a closer analysis of the tableau method reveals that it only uses polynomial space. Semantic tableaux lie in between model theory and proof theory, and they fit well with filtration. Indeed, there exist versions of filtration [39, 42, 44] that read-off counter-models by tableau-like recipes. A further comparison of tableau methods and filtration methods is beyond the scope of this article.

Another basic feature of modal languages, and the reason behind many of their practical applications, is their low model-checking complexity. Unlike first-order logic, whose model-checking problem is *Pspace*-complete, model checking for the basic modal language is known to be in low *Ptime*. Interestingly, the efficient algorithm which shows this computes truth values for the given formula and all its subformulas in all worlds of the given model. The reason why this suffices, suggests analogies with filtration, but again, we forego a further discussion.[10]

All this suggests that filtration can be constructivized even further, as a way of getting at the minimal complexity of basic tasks associated with modal logics.

3.4 Conclusion

The method of filtration can also be used in a proof-theoretic setting, and then it adds additional constructive insights into axiomatization.

3.5 Digression: Algebraic versions

This section was about merging semantic and proof-theoretic aspects of filtration. A standard setting for such a two-sided perspective is the algebraic approach to modal logic. We briefly discuss an algebraic version of filtration which goes back to [58, 59]. For recent accounts, see [18, 32, 46] – while we will follow [18].

Recall that a *modal algebra* is a pair $\mathfrak{A} = (A, \Diamond)$, where A is a Boolean algebra and $\Diamond$ is a unary function on A that commutes with finite joins. As usual, the dual operator $\Box$ is defined as $\neg\Diamond\neg$. A *modal homomorphism* between two modal algebras is a Boolean homomorphism h satisfying $h(\Diamond a) = \Diamond h(a)$. Let MA be the category of modal algebras and modal homomorphisms.

Definition 3.9 *Let $\mathfrak{A} = (A, \Diamond)$ and $\mathfrak{B} = (B, \Diamond)$ be modal algebras and let $h : A \to B$ be a Boolean homomorphism.*

[10] The low complexity also implies that computing filtrated models can be done efficiently.

1. *A map h is a* stable homomorphism *provided* $\Diamond h(a) \leq h(\Diamond a)$ *for each* $a \in A$.
2. *A map h satisfies* the closed domain condition (CDC) for $D \subseteq A$ *if* $h(\Diamond a) \leq \Diamond h(a)$ *for each* $a \in D$.

It is easy to see that $h : A \to B$ is stable iff $h(\Box a) \leq \Box h(a)$ for each $a \in A$. A valuation on a modal algebra $(A, \Diamond)$ is a map V from propositional variables to A. Each valuation can be extended to a map from the set of all formulas into A in a standard way. We write $(A, \Diamond, V) \models \varphi$ if $V(\varphi) = 1$.

Definition 3.10 *Let* $\mathfrak{A} = (A, \Diamond)$ *be a modal algebra, V a valuation on A, and* Σ *a set of formulas closed under subformulas. Let* A^f *be the Boolean subalgebra of A generated by* $V(\Sigma) \subseteq A$ *and let* $D = \{V(\varphi) : \Diamond\varphi \in \Sigma\}$. *Let* $\Diamond^f$ *be a modal operator and* V^f *a valuation on* A^f. *Then* $(A^f, \Diamond^f, V^f)$ *is called an* algebraic filtration *of* $(A, \Diamond, V)$ *through* Σ *if* $V^f(p) = V(p)$ *for each* $p \in \Sigma$ *and*

1. *the inclusion* $(A^f, \Diamond^f) \rightarrowtail (A, \Diamond)$ *is a stable homomorphism, i.e.,* $\Diamond^f a \leq \Diamond a$, *for each* $a \in A$;
2. *the inclusion* $(A^f, \Diamond^f) \rightarrowtail (A, \Diamond)$ *satisfies* (CDC) *for D, i.e.,* $\Diamond a \leq \Diamond^f a$, *for each* $a \in D$.

Using a duality of modal algebras one can show that algebraic filtration is dual to standard model-theoretic filtration, [54, 55]. In fact, condition (1) is an algebraic analogue of condition (F1) of filtration and condition (2) is the analogue of (F2). The next theorem is an algebraic analogue of the basic filtration theorem.

Theorem 3.11 (Algebraic filtration theorem) *Let* $(A, \Diamond)$ *be a modal algebra, V a valuation on A and* Σ *a set of formulas closed under subformulas. Let* $(A^f, \Diamond^f, V^f)$ *be a filtration of* $(A, \Diamond, V)$ *through* Σ. *Then for every* $\varphi \in \Sigma$ *we have*

$$V(\varphi) = V^f(\varphi).$$

For recent general algebraic accounts of filtration, see [18, 32, 46].

Weak filtrations introduced in Section 2.4 then correspond to algebras $(A^f, \Diamond^f, V^f)$ such that the inclusion $(A^f, \Diamond^f) \rightarrowtail (A, \Diamond)$ satisfies $\Diamond a = \Diamond^f a$, for each $a \in D$. Thus, algebraic filtration can be seen as a way of 'twisting' the standard denotation of the standard modal operator in such a way that we do not get full, potentially infinite, subalgebras, but finite subalgebras whose modal operator is just enough like the standard one to satisfy the relevant formulas.

Remark 3.12 The algebraic approach links naturally to later topics in this article. In particular, it lends itself to the dynamics of model update discussed in Sections 4 and following. Given a set of formulas Σ and a modal algebra with a valuation $(A, \Diamond, V)$ we can define a dynamic operation $[\Sigma]\varphi$ by putting $(A, \Diamond, V) \models [\Sigma]\varphi$ iff $(A^f, \Diamond^f, V^f) \models \varphi$, where $(A^f, \Diamond^f, V^f)$ is a filtration of $(A, \Diamond, V)$ via Σ. This suggest an algebraic version of filtration dynamics. From this perspective 'filtration as abstraction' is encoded in the fact that we are generating A^f from $V(\Sigma)$. Thus, algebraically, generating a subalgebra corresponds to filtering information through Σ.

Many further dynamic operations make sense on algebras. For example, given $(A,\Diamond,V)$ and a formula ψ, let $(A,\Diamond,V) \models [!\psi]\varphi$ iff $(A_\psi,\Diamond_\psi,V_\psi) \models \varphi$, where $(A_\psi,\Diamond_\psi,V_\psi)$ is obtained from $(A,\Diamond,V)$ by relativizing its domain to the element $V(\psi)$. This is an algebraic analogue of public announcement, a notion whose logic PAL occurs in Section 5.1.[11] An algebraic analysis in this spirit is given in [53].

4 The modal logic of filtration

4.1 Filtration modalities

The basic theory of filtration as explained in Section 2 itself has a modal flavor. Filtration is a form of model change, and we can describe this change by introducing a new modal operator for it. This suggests an extension of the basic modal language with modalities $[\backslash\Sigma]\varphi$ saying that φ is true at the image of the current world in the result of filtrating the current model with respect to the set of formulas Σ.

If we fix the standard filtration, then this extension of the basic language of modal logic already adds expressive power. The following is easy to see.

Example 4.1 *The formula $[\backslash\{\top\}]\Diamond\top$ is equivalent to $\exists x\exists yRxy$. This is easy to see since the filtrated model for $\{\top\}$ contains just one world. But $\exists x\exists yRxy$ is not definable in the basic modal language since it is not invariant under bisimulation.*

The latter statement is definable with a universal modality in the language, as $\neg U\Box\bot$. Conversely, the universal modality is not definable in terms of the filtration modality, see [51, Chapter 8]. However, as we shall see soon, the basic modal language with added filtration modalities is closely related to the basic modal language with a universal modality added.

But we can go much further in this style of analysis. As we have seen, there are many variants of filtration, where relations in the filtrated model can be defined in various ways. One general format for defining new relations in a modal setting is the 'program format' of van Benthem & Liu [14], where new relations are given by programs in the language of propositional dynamic logic PDL, that we have used on several occasions already. We will see further examples in a moment.

The results that follow can be found in Baltag et al. [4] and Ilin [51]. We consider a language for the logic of filtration [4, 51] [12] which arises by extending the language of propositional dynamic logic PDL with abstraction modalities $[\overrightarrow{\pi}/\Sigma]\varphi$.

---

[11] One can also consider more purely algebraic operations. E.g., let $f : A \to A$ be any map, with A_f the set of fixed-points of f, and set $A \models [f]\varphi$ iff $A_f \models \varphi$. Here A could be a Heyting algebra and f a nucleus on it, say, the double negation $\neg\neg$. Then A_f is a Boolean algebra of regular elements of A, cf. [20]. Algebraic machinery of this sort fits well with various later topics in this article, but we must leave the exploration of algebraic perspectives to another occasion.

[12] The cited works call this general system a 'logic of abstraction', emphasizing the fact that filtration stands for a very general procedure.

Definition 4.2 *The* dynamic language $\mathsf{PDL}_{[\overrightarrow{\pi}/\Sigma]}$ *is defined by the grammar:*

$$\pi := r \mid ?\psi \mid 1 \mid \pi;\pi \mid \pi \cup \pi \mid \pi^* \quad \textit{and}$$
$$\varphi := p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \langle\pi\rangle\varphi \mid [\overrightarrow{\pi}/\Sigma]\varphi,$$

where r is a basic program in Π_0, 1 is the universal relation, $\psi \in \mathsf{PDL}_{[\overrightarrow{\pi}/\Sigma]}$, $\overrightarrow{\pi} = (\pi_r)_{r\in\Pi_0}$ is a sequence of PDL*-programs, and Σ is a* finite[13] *subset of* PDL.

Next we define the relevant (multi-relational) models.

Definition 4.3 (Quotient model) *Let $\mathfrak{M} = (W, (R_a)_{a\in\Pi_0}, V)$ be a relational model. For any finite $\Sigma \subseteq \mathsf{PDL}$ and any sequence $\overrightarrow{\pi} = (\pi_a)_{a\in\Pi_0}$ of programs, the* quotient model $\mathfrak{M}^{\overrightarrow{\pi}}_{\Sigma}$, *is $\mathfrak{M}^{\overrightarrow{\pi}}_{\Sigma} = (W_\Sigma, (R^{\pi_a}_{\Sigma})_{a\in\Pi_0}, V_\Sigma)$, where*

- $W_\Sigma := \{[w]_\Sigma \mid w \in W\}$,
- $V_\Sigma(p) := \{[w]_\Sigma \mid$ *there is* $w' \sim_\Sigma w$ *with* $w' \in V(p)\}$, *and*
- *For each $a \in \Pi_0$,*

$$[w]_\Sigma R^{\pi_a}_{\Sigma} [v]_\Sigma \textit{ iff there is } w' \sim_\Sigma w \textit{ and } v' \sim_\Sigma v \textit{ with } w' R_{\pi_a} v',$$

where the relation $\sim_\Sigma$ is the equivalence relation induced by Σ defined just as in Section 2.1 and R_π denotes the relation induced by the program π.

Recall the notions of smallest, largest, transitive, and smallest transitive filtration in Section 2.1, for which we use the abbreviations s, l, t, st, respectively. To fit these into our program format, for each $f \in \{s, l, t, st\}$, we define a program π_f in the language PDL_{-*} (*-free PDL) whose quotient models coincide with f-filtrations. Let Σ be a finite set of formulas in the language $\mathscr{L}_E$. For $\Psi \subseteq \Sigma$, we set

- $\Psi_\Diamond := \bigwedge_{\Diamond\varphi\in\Sigma,\varphi\in\Psi} \Diamond\varphi$,
- $\Psi_{\Diamond,\vee} := \bigwedge_{\Diamond\varphi\in\Sigma,\varphi\in\Psi} (\Diamond\varphi \vee \varphi)$,
- $\neg\Psi := \{\neg\varphi \mid \varphi \in \Psi\}$.

while we also let $\widehat{\Psi} = \bigwedge\Psi \wedge \bigwedge\neg(\Sigma\setminus\Psi)$,

Then we define the following programs:

$$\pi_s = r, \quad \pi_l = \bigcup_{\Psi\subseteq\Sigma} (?\Psi_\Diamond; 1; ?\widehat{\Psi}), \quad \pi_t = \bigcup_{\Psi\subseteq\Sigma} (?\Psi_{\Diamond,\vee}; 1; ?\widehat{\Psi}).$$

Now let $\pi_\Sigma = \bigcup_{\Psi\subseteq\Sigma} (?\widehat{\Psi}; 1; ?\widehat{\Psi})$, and for $k \in \mathbb{N}$, let $\pi_1 = r$ and $\pi_{k+1} = r; \pi_\Sigma; \pi_k$, while for the smallest transitive filtration, we set

$$\pi_{st} = \bigcup_{1\leq k\leq 2^{|\Sigma|}} \pi_k.$$

[13] The finiteness is essential for obtaining reduction axioms for the corresponding dynamic logic.

Lemma 4.4 *Let $f \in \{s, l, t, st\}$, Σ a finite subformula closed set, and $\mathfrak{M}$ any model. The quotient model $\mathfrak{M}_\Sigma$ w.r.t. the program π_f matches the f-filtration of $\mathfrak{M}$ via Σ.*

For a proof we refer to [4] and [51, Chapter 8].

Definition 4.5 (Semantics for $\mathsf{PDL}_{[\overrightarrow{\pi}/\Sigma]}$) *Let $\mathfrak{M} = (W, (R_a)_{a\in\Pi_0}, V)$ be a relational model and w in W. The truth of $\mathsf{PDL}_{[\overrightarrow{\pi}/\Sigma]}$-formulas is defined recursively as for PDL, now with the additional clause:*

$$\mathfrak{M}, w \models [\overrightarrow{\pi}/\Sigma]\varphi \;\text{ iff }\; \mathfrak{M}^{\overrightarrow{\pi}}_{\Sigma}, [w]_\Sigma \models \varphi$$

where the model $\mathfrak{M}^{\overrightarrow{\pi}}_{\Sigma}$ is as described in Definition 4.3.

4.2 Completeness theorem

One can axiomatize the theory of filtration in the above sense as a quotient logic.

For a start, for every formula $\chi \in \mathsf{PDL}_{[\overrightarrow{\pi}/\Sigma]}$ and finite $\Sigma \subseteq \mathsf{PDL}$ we set:

$$\langle\sim_\Sigma\rangle\chi := \bigvee_{\Psi\subseteq\Sigma}\left(\widehat{\Psi} \wedge \langle 1\rangle\left(\widehat{\Psi} \wedge \chi\right)\right).$$

The following lemma shows that the modality $\langle\sim_\Sigma\rangle$ is in fact the standard diamond modality for the binary relation $\sim_\Sigma$.

Lemma 4.6 *For a modal model $\mathfrak{M}$ and world $x \in \mathfrak{M}$, we have*

$$\mathfrak{M}, x \models \langle\sim_\Sigma\rangle\chi \text{ iff there is } x' \sim_\Sigma x \text{ with } \mathfrak{M}, x' \models \chi.$$

Next, the logic QPDL is determined by the axioms and rules in Table 2.

(PDL)	Axiom-schemes and rules of PDL (see Table 1)
(Ax-$\langle 1\rangle$)	S5-axioms for $\langle 1\rangle$, $\langle\pi\rangle\varphi \to \langle 1\rangle$
(Ax-$\mathsf{K}_{[\overrightarrow{\pi}/\Sigma]}$)	$[\overrightarrow{\pi}/\Sigma](\varphi \to \psi) \to ([\overrightarrow{\pi}/\Sigma]\varphi \to [\overrightarrow{\pi}/\Sigma]\psi)$
(Ax-p)	$[\overrightarrow{\pi}/\Sigma]\, p \leftrightarrow \langle\sim_\Sigma\rangle p$
(Ax-$\neg$)	$[\overrightarrow{\pi}/\Sigma]\neg\varphi \leftrightarrow \neg[\overrightarrow{\pi}/\Sigma]\varphi$
(Ax-$\wedge$)	$[\overrightarrow{\pi}/\Sigma](\varphi \wedge \psi) \leftrightarrow [\overrightarrow{\pi}/\Sigma]\varphi \wedge [\overrightarrow{\pi}/\Sigma]\psi$
(Ax-$\langle 1\rangle$)	$[\overrightarrow{\pi}/\Sigma]\langle 1\rangle\varphi \leftrightarrow \langle 1\rangle[\overrightarrow{\pi}/\Sigma]\varphi$
(Ax-$\langle r\rangle$)	$[\overrightarrow{\pi}/\Sigma]\langle r\rangle\varphi \leftrightarrow \langle\sim_\Sigma\rangle\langle\pi_r\rangle[\overrightarrow{\pi}/\Sigma]\varphi$ for all $r \in \Pi_0$
(Ax-$*$)	$[\overrightarrow{\pi}/\Sigma]\langle\alpha^*\rangle\varphi \leftrightarrow [\overrightarrow{\pi}/\Sigma]\bigvee_{0\le n\le 2^{\lvert\Sigma\rvert}}\langle\alpha\rangle^n\varphi$
(Nec$_{[\overrightarrow{\pi}/\Sigma]}$)	From φ infer $[\overrightarrow{\pi}/\Sigma]\varphi$

Table 2: The logic QPDL

We can now formulate the main result.

Theorem 4.7 (Soundness and completeness) *The logic* QPDL *is sound and complete with respect to the semantics of the modal language of filtration.*

The key observation in the proof of this result is that the stated reduction axioms enable us to show that every formula in $\mathsf{PDL}_{[\overrightarrow{\pi}/\Sigma]}$ is provably equivalent in the system QPDL to a formula in the language PDL.

Theorem 4.8 (Expressivity) *For every* $\varphi \in \mathsf{PDL}_{[\overrightarrow{\pi}/\Sigma]}$ *there is a* $\psi \in \mathsf{PDL}$ *such that* $\vdash_{\mathsf{QPDL}} \varphi \leftrightarrow \psi$.

The latter result is proved by repeating the following procedure. Start with any innermost occurrence of a filtration modality in φ, with no such modalities in its scope, and use the axioms displayed to push this occurrence inside until it reaches atomic formulas. Using our axiom for this base case, each of these last modalized subformulas can be replaced by a formula without filtration modalities.[14]

4.3 Discussion

There are several things to note about the results obtained here, since the above analysis has several interesting consequences.

Inside and outside the filtration set. First, we note that the original filtration theorem is a special case. The logic QPDL is complete with respect to formulas $[\overrightarrow{\pi}/\Sigma]\varphi$ where φ is arbitrary, and need not occur inside the set Σ. In the special case that $\varphi \in \Sigma$, however, we can get something much stronger.[15] Suppose that all relations are treated via the coarsest filtration, then the standard filtration equivalence

$$[\backslash\Sigma]\varphi \leftrightarrow \varphi$$

will be provable. This can be shown by induction on φ, where the case for the existential modality $\Diamond$ involves a non-trivial appeal to the principles of QPDL.[16]

But our results also apply to well-known further topics in classical modal logic.

Subframe logics and stable logics. Fine's special modal logics mentioned in Section 2.1 reappear in this setting. Call a modal logic L *closed under updates* if, given any model $\mathfrak{M}$ based on a frame for L, for each formula φ, the updated model $\mathfrak{M}|\varphi$, i.e., the relativization of $\mathfrak{M}$ to φ, is also based on a frame for L. It is easy to verify that a modal logic L is closed under updates if and only if it is a subframe logic.

[14] The precise version of this procedure requires a non-trivial termination argument in terms of a syntactic measure on formulas whose details can be found in [51, Chapter 8].

[15] Here and in what follows, we suppress the explicit program notation π of QPDL, since the definition of the relations is fixed for the standard filtration.

[16] The reader may find of interest to trace the details of this proof, and see how it literally matches the usual argument for the filtration theorem stated in the meta-language.

In a similar way we can define a modal logic L to be *closed under abstraction* if given a model $\mathfrak{M}$ based on a frame of L, for each set of formulas Σ, the updated model $\mathfrak{M}^{f,\Sigma}$ is also based on a frame of L. These are the logics for which it is meaningful to consider the abstraction/filtration dynamics. It follows from [19, 51] that stable logics (in the language with the universal modality), mentioned in Section 2.1, are exactly the logics that are closed under abstraction.

Next we turn to some less standard perspectives.

Schematic validity. A third interesting point is that there are natural schematically valid principles (all their substitution instances are universally valid) concerning filtration that do not occur in the above axiomatization. One example is *Idempotence*:

$$[\backslash\Sigma]\varphi \leftrightarrow [\backslash\Sigma]\,[\backslash\Sigma]\varphi$$

Not all principles of our logic QPDL are schematically valid: e.g., it is easy to see that the above Axiom (Ax-p) is not. This leaves us with an interesting

Open problem 4.9 *Can the schematic validities of the modal logic of filtration be axiomatized, and if so, how?*

Fixed-point logic. Another interesting open problem concerns the cavalier treatment of iteration via brute enumeration in the logic QPDL, which used the finiteness of filtrated models with a bound given by the set Σ. There is an obvious problem of how to generalize this to arbitrary settings with infinite sets Σ, in itself a natural generalization of the filtration setting.[17]

Limits to axiomatizing meta-theory. Despite the above success, how much meta-model theory can one axiomatize in modal logic? For the case of first-order logic, a striking mismatch was pointed out in [57]: the elementary meta-theory of first-order predicate logic encodes True Arithmetic, and hence it is non-arithmetical, and hence undecidable and non-axiomatizable. A similar mismatch may well occur for modal logic. For instance, the Lindström-style analysis of modal logic in [16] shows how the crucial notion of a bisimulation is not itself modally definable.[18]

4.4 Conclusion

This section has shown that 'modalizing filtration' is a feasible strategy. The basic theory of this method is so simple that it can itself be absorbed into modal logic. This finding points the way to a more ambitious program of formalizing 'quotient dynamics' that we will return to in Section 6. It also raises a broader foundational

[17] A modal logic with 'ceteris paribus' riders expressed as arbitrary sets of formulas whose truth values are to be kept constant when comparing worlds is studied in [13]. Its connection to the above logic of filtration remains to be understood.

[18] The general modal logic of bisimulation over universes of relational models seems undecidable, since bisimulation imposes a form of commutation that can encode grid structure, [25].

issue, namely, to which extent a modal encoding of the meta-theory of modal logic is possible, something that will be considered briefly in Section 8.

5 Filtration and information update

In the preceding section, we saw how, in a semantic perspective, filtration is a natural form of model change that might be described as 'definable quotienting'. But many other definable transformations of models are studied in modal logic today, in particular, in the tradition of *dynamic-epistemic logic* where information update is represented by model change, [9]. In fact, the above axiomatization of modal filtration logic used a well-known dynamic-epistemic technique, describing the behavior of the additional dynamic modalities in terms of 'recursion axioms'.

These two ways of studying model change do not live in isolation, it makes sense to compare and combine them, and in this section, we will explore how this works.[19]

5.1 Public announcement logic

The simplest information update arises when new information comes in as a true proposition φ, shrinking the current range of options for what the actual world is like. Technically, this is relativization to a definable submodel: a current pointed epistemic model $\mathfrak{M} = (W,R,V),s$, where s is the actual world, is transformed into

$$\mathfrak{M}|\varphi, s$$

that is, the submodel consisting only of those points that satisfied φ in $\mathfrak{M}$. This update is a partial function: the transformation is only defined when $\mathfrak{M},s \models \varphi$: that is, the new information is true. These actions are often called 'public announcements', though they can also result from observation, or any signal that the agent decides to treat as completely reliable.

The modal language of *public announcement logic* PAL to be used here consists of the standard monomodal language plus action terms $!\varphi$ for any formula of the language, and dynamic modalities $[!\varphi]\psi$ for announcements interpreted as follows:

$$\mathfrak{M},s \models [!\varphi]\psi \quad \text{iff} \quad (\text{if } \mathfrak{M},s \models \varphi, \text{ then } \mathfrak{M}|\varphi,s \models \psi)$$

The existential modality $\langle !\varphi\rangle\psi$ makes the same assertion about the relativized submodel, but now in conjunction with $\mathfrak{M},s \models \varphi$. We will use this version occasionally.

A complete proof system for PAL consists of the axioms and rules for the minimal normal modal logic, now for the extended language, together with the four *recursion axioms* displayed in Table 3.

[19] This section is discursive, putting a topic on the map, rather than presenting deep new results.

Reduction laws:	
(R_{atom})	$[!\theta]p \leftrightarrow (\theta \to p)$
($R_{\neg}$)	$[!\theta]\neg\psi \leftrightarrow (\theta \to \neg[!\theta]\psi)$
($R_{\wedge}$)	$[!\theta](\psi \wedge \varphi) \leftrightarrow ([!\theta]\psi \wedge [!\theta]\varphi)$
($R_{\Diamond}$)	$[!\theta]\Diamond\psi \leftrightarrow (\theta \to \Diamond[!\theta]\psi)$
(R_U)	$[!\theta]U\varphi \leftrightarrow (\theta \to U[!\theta]\varphi)$
($R!!_{[!]}$)	$[!\theta][!\rho]\chi \leftrightarrow [!(\theta \wedge [!\theta]\rho)]\chi$
($Nec_!$)	from φ, infer $[!\theta]\varphi$.

Table 3: Axioms and rules of PAL

Theorem 5.1 *The displayed proof system for* PAL *is sound and complete.*

Proof. The method is like that used in Section 4 for the modal logic of filtration. Working iteratively, innermost occurrences of dynamic modalities can be pushed inside using the recursion axioms and the proof rule of Replacement of Equivalents, until they attach to atoms, where the dynamic modality can be removed by the axiom for atoms. For the remaining valid formulas without dynamic modalities, there must be a proof given the completeness of Basic Modal Logic.

This reduction method also establishes decidability, since validity in PAL can now be decided by applying the transformation and then testing for validity of the resulting formula in Basic Modal Logic.[20] [21]

While this reduction method is convenient, there are also more general completeness proofs that work in more standard Henkin-style. These also work for variations of the system that no longer support recursion axioms of the above kind.

There are many further dynamic logics of this kind, covering updates with private as opposed to public information, belief revision, preference change, issue change, and in fact, any activity of rational agents that can be affected by new information. In this broader setting, more common than the domain change in the above is a transformation of 'definable relation change' on an unchanged domain. In Section 4, we already used the 'program format' that is often used in dynamic-epistemic logic for definable relation change.

[20] This method involves potential exponential blow-up, and does not establish the complexity of satisfiability for PAL, which is in fact *Pspace*-complete.

[21] Incidentally, a direct semantic filtration argument for the decidability of PAL is not easy to give. Could the results that follow provide a principled solution?

5.2 Quotient dynamics and update dynamics

In handling information, both perspectives make sense: update as changing a current range of options, and quotienting as redescribing that range for some special purpose at hand. What happens when we combine the two?

In what follows, unless explicitly mentioned, we assume that the sets of formulas used for filtration are closed under subformulas.[22]

Fact 5.2 *The following commutation principle holds:* $!\varphi\ ;\ \backslash\Sigma\ \cong\ \backslash\Sigma_\varphi\ ;\ !\varphi$, *with* $\Sigma_\varphi = \{\langle\varphi\rangle\alpha \mid \alpha \in \Sigma\}$ *and with ; for relational composition. That is, when applied to the same model* $\mathfrak{M}, s$, *the two operator sequences yield isomorphic models.*

$$\begin{tikzcd}
\mathfrak{M} \arrow[rr, "!\varphi"] \arrow[d] & & \mathfrak{M} \mid \varphi \arrow[d] \\
\mathfrak{M}^{f,\Sigma_\varphi} \arrow[r, "!\varphi"'] & (\mathfrak{M}^{f,\Sigma_\varphi}) \mid \varphi & \cong\ (\mathfrak{M} \mid \varphi)^{f,\Sigma}
\end{tikzcd}$$

Proof. Here is the key observation in showing that the two routes to the bottom-right corner of the diagram are the same. In the model $\mathfrak{M}|\varphi$ after update with φ, points s,t that are Σ-equivalent are mapped onto the same point in $(\mathfrak{M}|\varphi)^{f,\Sigma}$. But by the semantics of PAL, these same points s,t were equivalent in $\mathfrak{M}$ with respect to the formulas $\langle!\varphi\rangle\alpha$ for each $\alpha \in \Sigma$. We can also assume that $\langle!\varphi\rangle\top$ is among the relevant formulas [23]: in $\mathfrak{M}$, this formula marks the points satisfying φ. By the closure under subformulas, and the basic filtration theorem, the images of these points are exactly the points that satisfy φ in the filtrated model $\mathfrak{M}^{f,\Sigma_\varphi}$. And the latter are the points that are preserved in the bottom-right model $(\mathfrak{M}^{f,\Sigma_\varphi}) \mid \varphi$.

The final part of the proof shows that the bijection between $(\mathfrak{M} \mid \varphi)^{f,\Sigma}$ and $(\mathfrak{M}^{f,\Sigma_\varphi}) \mid \varphi$ found in this way is an isomorphism. This requires a routine argument chasing the diagram and using two main features: (i) the relational definition of the coarsest filtration, and (ii) the fact that updates $!\varphi$ lead to submodels, structures that inherit the old ordering among their points. We suppress the details here. □

This result says that updates at the richer level of the original models can be mimicked faithfully at the coarser level of filtrated models.

What cannot work, however, is finding a commuting diagram as above, but now with the filtration set Σ kept fixed.

[22] Even so, we will allow arbitrary announcements of new information, without considering what happens when we restrict new information to formulas inside or close to the filtration set.

[23] This follows from subformula closure, or just by putting $\top \in \Sigma$.

Example 5.3 Consider a model $\mathfrak{M}$ with two points s,t with just one R-arrow, running from s to t. Now filtrate with respect to the set $\Sigma = \{\top\}$. The result is one-point reflexive model. Next, consider the public announcement $!\Diamond\top$ on $\mathfrak{M}$: the result is a one-point irreflexive model. Filtrating the latter model with $\Sigma = \{\top\}$ just yields the same model. But there is no public announcement that takes a one-point reflexive model into a one-point irreflexive one.[24]

Background: Tracking. The general background for the preceding observations is the subject of 'tracking' information at various levels of structure, in the sense of validating commuting diagrams of the sort discussed above, cf. [10]. Tracking is important, since it connects and compares different approaches to information structure, for instance, intensional or more hyper-intensional. We refer to the cited paper for further concrete examples, and general results on when, given some projection functor from the finer to the coarser level of models, model transformations at richer levels are trackable by matching transformations at the coarser level. It is also shown there that, going in the opposite direction from coarser to finer, tracking is always possible in principle, but see below.

However, interestingly, our discussion involved a new feature: even when tracking in the strict sense of [10] is impossible, it may become possible by also varying the projection map, as we did above by adjusting the sets one filtrates with.[25]

Even so, issues remain in the present setting. Consider the inverse of the above scenario: $\backslash\Sigma$; $!\varphi$ gives an update on the coarser filtrated model, and we ask for a matching update on the richer original model. In principle, as we said, this is always possible: one just defines an abstract update function on the points of $\mathfrak{M}$ by taking those points s for which $s\backslash_\Sigma$ satisfies φ. But the more relevant issue is whether there is a *definable* matching update, produced by modifying the formula that is publicly announced. Here a natural candidate would be the following isomorphism:

$$\backslash\Sigma\ ;\ !\varphi \ \cong\ ![\backslash\Sigma]\varphi\ ; \backslash\Sigma.$$

But this principle is not valid. To refute it, consider a modal model $\mathfrak{M}$ with domain $W = \{s,t,u\}$, relation $R = \{(s,s),(t,t),(t,u),(u,t),(u,u)\}$, and valuation $V(p) = \{s,t\}$. Let the filtration set be $\Sigma = \{p, \Diamond\neg p\}$. The filtrated model is isomorphic to the model $\mathfrak{M}$ itself. If we then update this model with $!p$, we arrive at the submodel $\mathfrak{M}|p$ whose domain is $\{s,t\}$. But if we first update the initial model $\mathfrak{M}$ with $!p$ and only then filtrate that model with respect to Σ, a one-point model results where p holds, clearly not isomorphic to the two-element model obtained before.

We believe there are ways around this problem, but have not yet found one.[26]

---

[24] This argument no longer works if we allow a richer set of dynamic-epistemic model transformations such as changing relations. We leave an analysis of this broader setting to further work.

[25] Notice that, since formulas can change truth values under model change, our Σ-dependent notion of filtration makes our functor context-dependent in a way not envisaged in [10].

[26] One option might be to relax the notion of isomorphism in our diagrams to modal bisimulation. Moreover, one should note that solutions need not be unique. Another approach would keep the update the same, but change the filtration set, as in: $\backslash\Sigma\ ;\ !\varphi = !\varphi\ ;\ \backslash\Sigma^\varphi$ with $\Sigma^\varphi = \{\alpha | \langle\varphi\rangle\alpha \in \Sigma\}$.

5.3 Merging filtration logic with dynamic logic

The preceding results can be formulated in terms of a relative completeness result, in the combined language of filtration dynamics and public announcement.

Let UFL be the proof system resulting from combining the axioms and rules of the systems PAL and QPDL, together with the following axiom:

$$\langle !\varphi\rangle\, [\backslash\Sigma\,]\, \psi \leftrightarrow [\backslash\Sigma^{\varphi}\,]\, \langle !\varphi\rangle \psi,$$

with $\Sigma^{\varphi} = \{\langle\varphi\rangle\alpha \,|\, \alpha \in \Sigma\}$ as above.

Theorem 5.4 *The logic* UFL *is complete relative to its static base logic.*

The proof is a dynamic-epistemic-style reduction argument, where the new recursion axioms helps push announcement modalities inside across filtration modalities.

Note that this theorem does not immediately settle the decidability of UFL, since we end up in a modal filtration logic allowing different vocabularies inside its modalities. We will return to this setting in the next section.

5.4 Conclusion

This section may not contain striking new results, but it has put a new general issue on the map that makes sense in practice: interleaving filtration dynamics and modal update dynamics. Of course, we have considered just a few scenarios. It would also be natural, for instance, to let new announcements raise new issues that would change the current filtration set. We leave such scenarios to further investigation.

6 Issues and questions

6.1 Issues and questions

Filtration may be said to answer the question why a given formula φ is true at a given point s in a model $\mathfrak{M}$. In other words, φ is the issue to be resolved, and we do this by looking at a number of closely connected issues. This setting is very similar to one known from the semantics of questions [47]. In general, questions raise issues, and at any stage, the current issue can be seen as a *partition* of the current model. Resolving the issue means finding out in which partition cell we are, for instance, through information updates of the sort discussed in the preceding section.[27]

[27] Inquisitive logics [31], use generalized partitions, but this feature is orthogonal to our discussion.

Admittedly, the analogy is not perfect. In the semantics of questions, one usually keeps the whole domain of worlds unchanged, since these worlds represent the total relevant informational setting. In filtration, one reduces that model to a bare minimum with as few worlds as possible needed for resolving the current issue. Even so, it makes sense to look into some connections between logics of questions, logics of filtration, and the dynamics of refining partitions.

6.2 Modal logic of issues and filtration

A modal logic for representing information and issues is presented in [15]. It works on S5-models where one relation $\sim$ stands for an information partition, or an S5-style epistemic equivalence relation, and another equivalence relation $\approx$, whose equivalence classes represent the different complete answers to the current question, or more generally, the different ways of resolving the current issue. There are two matching standard modalities:

- a standard modality $\Box\varphi$ stands for the agent's information or knowledge at a world: it says that φ is true in all $\sim$-accessible worlds,
- a new 'issue modality' $Q\varphi$ that stands for truth of φ in all $\approx$-accessible worlds, describing the actual issue cell that inquiry should, ideally, get us to.[28]

This logic can be axiomatized in a standard manner, as a modal logic with two S5 modalities. On this basis, one can reformulate the earlier logic of filtration in Section 4 in more abstract, and perhaps more elegant, terms as a modal logic of quotient formation. The following result is taken from [4].

Models $\mathfrak{M} = (W, (R_a)_{a\in\Pi_0}, Q, V)$ consist of a relational model $(W, (R_a)_{a\in\Pi_0}, V)$ and an equivalence relation Q on W. We then define a language $\mathsf{PDL}_{Q,\overrightarrow{\pi}/Q}$ as:

$$\pi := r \mid Q \mid ?\psi \mid 1 \mid \pi;\pi \mid \pi\cup\pi \mid \pi^*, \text{ and}$$
$$\varphi := p \mid \neg\varphi \mid \varphi\wedge\varphi \mid \langle\pi\rangle\varphi \mid [\overrightarrow{\pi}/Q]\varphi,$$

where r is an element of the set of the basic programs Π_0 and ψ is a formula of PDL (the language $\mathsf{PDL}_{Q,\overrightarrow{\pi}/Q}$ without occurrences of $[\overrightarrow{\pi}/Q]\varphi$).

Let $\mathfrak{M} = (W, (R_a)_{a\in\Pi_0}, Q, V)$ be a model. We let $[w]$ denote the equivalence class of w with respect to Q. For a sequence of programs $\overrightarrow{\pi}$, we define a model $\mathfrak{M}_Q^{\overrightarrow{\pi}} := (W_Q, (R_Q^{\pi_a})_{a\in\Pi_0}, \mathsf{Id}, V_Q)$, where $W_Q := \{[w] \mid w\in W\}$, $V_Q(p) := \{[w] \mid w\in V(p)\}$, Id denotes the identity relation, and

$$[w]R_Q^{\pi_a}[v] \text{ iff there is } w'Qw \text{ and there is } v'Qv \text{ such that } w'R_{\pi_a}v'.$$

The crucial step in the semantics is:

[28] Typically, an agent will not know which issue cell she is in, until the issue has been resolved.

$$\mathfrak{M}, x \models [\overrightarrow{\pi}/Q]\varphi \;\text{ iff }\; \mathfrak{M}^{\overrightarrow{\pi}}_{Q}, [x] \models \varphi.$$

To obtain a convenient representation of the key axioms in our issue/quotient logic, we define functions $f_{Q,\overrightarrow{\pi}}$ on programs by the following clauses:

$$f_{\overrightarrow{\pi},Q}(Q) = ?\top,$$
$$f_{\overrightarrow{\pi},Q}(r) = Q;\overrightarrow{\pi},$$
$$f_{\overrightarrow{\pi},Q}(\alpha_1 \circ \alpha_2) = f_{\overrightarrow{\pi},Q}(\alpha_1) \circ f_{\overrightarrow{\pi};Q}(\alpha_2) \text{ for } \circ \in \{\cup, ;\}$$
$$f_{\overrightarrow{\pi},Q}(\pi^*) = \left(f_{\overrightarrow{\pi},Q}(\pi)\right)^*.$$

The recursion axioms matching these stipulations are listed in Table 4.

(PDL)	Axiom-schemes and rules of PDL
(Q)	S5-axioms and rules for Q
(Ax-p)	$[\overrightarrow{\pi}/Q]p \leftrightarrow \langle Q\rangle p$
(Ax-$\neg$)	$[\overrightarrow{\pi}/Q]\neg\varphi \leftrightarrow \neg[\overrightarrow{\pi}/Q]\varphi$
(Ax-$\wedge$)	$[\overrightarrow{\pi}/Q](\varphi \wedge \psi) \leftrightarrow [\overrightarrow{\pi}/Q]\varphi \wedge [\overrightarrow{\pi}/Q]\psi$
(Ax-$\langle\alpha\rangle$)	$[\overrightarrow{\pi}/Q]\langle\alpha\rangle\varphi \leftrightarrow \langle f_{Q,\overrightarrow{\pi}}(\alpha)\rangle[\overrightarrow{\pi}/Q]\varphi$
(Ax-$\langle Q\rangle$)	$[\overrightarrow{\pi}/Q]\langle Q\rangle\varphi \leftrightarrow [\overrightarrow{\pi}/Q]\varphi$
(DR-Nec)	From φ infer $[\overrightarrow{\pi}/Q]\varphi$

Table 4: The logic PDL_Q

It is easy to see that this logic is sound and complete for its intended semantics.[29]

6.3 *Issue dynamics*

But as in Section 5, a natural further topic arises. Issues are intrinsically ephemeral, they change over time, and so, they must be updated. This dynamics is the key theme of [15], and we state some relevant facts. A Yes/No question φ? is a typical instance of issue dynamics. Given a current issue relation $\approx$, an incoming new question introduces one more issue to be resolved, that may cut across existing issue cells. More precisely, the question triggers a *refinement* of the current issue relation $\approx$ to a new

[29] In Section 4, the analogue of the modality $\langle Q\rangle$ was *definable* in the language $\mathsf{PDL}_{[\overrightarrow{\pi}/\Sigma]}$ (cf. Lemma 4.6), and therefore, it was not needed in the syntax.

relation that now also cuts all links between φ-points and $\neg\varphi$-points, an operation that can be written as follows in our earlier PDL program format:

$$(?\varphi;\approx;?\varphi)\cup(?\neg\varphi;\approx;?\neg\varphi)$$

This definable update again satisfies recursion axioms in the earlier dynamic-epistemic format. For instance, here are two characteristic valid equivalences,

$$\langle ?\varphi\rangle\Box\psi \leftrightarrow \Box\langle ?\varphi\rangle\psi$$

$$\langle ?\varphi\rangle Q\psi \leftrightarrow (\varphi\wedge Q(\varphi\to\langle ?\varphi\rangle\psi))\vee(\neg\varphi\wedge Q(\neg\varphi\to\langle ?\varphi\rangle\psi))$$

The cited paper also studies other natural updates on current issues, including deleting agenda items. And there are also natural updates mixing the agents' information with the goal of their inquiry, such as taking the intersection of the epistemic partition and the issue partition.

Remark: partition algebra. The update triggered by a question can also be described differently, in terms of an algebra of partitions. Each formula φ induces a partition $\approx_\varphi$ of the domain into the set of all worlds satisfying φ and those satisfying $\neg\varphi$. The above update is nothing but the intersection of this relation with the current $\approx$, or in partition terms, their 'joint refinement'. Another natural operations on partitions is their joint coarsening. This is not the union of two equivalence relations: if a partition is needed, one needs to take the reflective transitive closure of the union.[30] Another option is just freely using the regular operations of PDL on partitions, including relational composition, now working with 'generalized partitions'.

It would be an interesting task to axiomatize the equational validities of partition algebra for the relations $\sim\backslash_\Sigma$ used in filtration. In Section 4 we observed that $\sim\backslash_\Sigma$ is idempotent, but many obvious validities fail. For instance, for different sets Σ,Θ, the composition $\sim\backslash_\Sigma$; $\sim\backslash_\Theta$is not commutative, or even associative. A counter-example to commutativity is easily extracted from Example 5.3 above. [31]

6.4 Filtration logic with varying vocabularies

The perspective of partition algebra suggests a different look at our earlier modal logic of filtration. Different partitions may come from different sets of formulas, so we get a question that was not really pursued in Section 4, namely, what are valid principles of a filtration logic with dynamic modalities for different vocabularies?

It is interesting to see that intuitions need not always be clear here, even for simple questions. For instance, it seems natural to expect a form of 'monotonicity': refining a partition should not change conclusions already changed. But it is immediate that this principle is not valid:

---

[30] For the partitions induced by filtrations, this union will often, trivially, be the universal relation.

[31] One difficulty is that, after filtration by a set Σ, filtration by an arbitrary other set Θ makes no sense, as the similarity type of the filtrated model has become that of the vocabulary of Σ only.

Example 6.1 The implication $[\backslash\Sigma]\varphi \to [\backslash(\Sigma \cup \Theta)]\varphi$ is not valid.

Consider our earlier model $\mathfrak{M}$ consisting of one arrow, that is, two points s, t such that Rst. Filtration with $\Sigma = \{\top\}$ compresses this to one reflexive point. So, at t in $\mathfrak{M}$, the formula $[\backslash\Sigma]\Diamond\top$ holds. But filtrating with the larger set $\{\top, \Diamond\top\}$ leaves the model the same, and hence at t in $\mathfrak{M}$, the formula $[\backslash\Sigma]\Diamond\top$ does not hold.

As can be seen in the preceding model, the converse implication $[\backslash(\Sigma \cup \Theta)]\varphi \to [\backslash\Sigma]\varphi$ is not valid either. Retreating to coarser filtrations, too, is a complex move.

Thus, filtrating with different sets raises intuitive challenges.[32] Perhaps the truly valid laws are different. For instance, does filtrating with two sets of formulas that both contain Σ make the same formulas from Σ true?

Example 6.2 The following composition law is valid:

$$[\backslash\Sigma][\backslash\Theta]\varphi \leftrightarrow [\backslash[\backslash\Sigma]\Theta]\varphi.$$

This equivalence analyzes the effect of two successive filtrations w.r.t. Θ and Σ, where the first changes the model so that truth values of some modal formulas can change. The same effect can be reached by filtrating once, now w.r.t. the single set of all more complex modal formulas of the form $[\backslash\Sigma]\theta$ with $\theta \in \Theta$.

The proof is a straightforward verification of two ways of describing an equivalence between worlds in the original model. In all, the preceding discussion, preliminary and inconclusive as it is, may have shown the interest in the following

Open problem 6.3 *Axiomatize modal filtration logic with varying vocabularies.*

Moreover, as a subsidiary question, one would want a complete dynamic logic over this with public announcement and natural notions of issue change.

6.5 Conclusion

Like the previous one, this section has mainly suggested a perspective, this time, that of merging logics of filtration with logics of issues, in both dynamic and static varieties. No major new technical result has been found, but we have seen how filtration combines well with current semantic trends toward adding what might be called focus or purpose to standard structures.

7 Exploring the boundaries

The filtration method has its boundaries: it cannot work universally, since not all logics are decidable. In particular, one expects its reach to stop at first-order logic. But more can be said, and we will explore the boundaries more precisely.

[32] Other non-valid principles include the putative equivalence of $\backslash\Sigma\ ;\ \backslash\Theta$ and $\backslash(\Sigma \cup \Theta)$.

7.1 Finite type models for first-order logic

Finding finite types, the hallmark of filtration, still works for a first-order language. Look at any first-order formula φ and take the finite set of its subformulas $Sub(\varphi)$. Any first-order model $\mathfrak{M}$ and variable assignment s induces a set

$$\text{Type}(\mathfrak{M}, s) = \{\psi \in Sub(\varphi) \mid \mathfrak{M} \models \psi\}.$$

In the same way one can define induced X-types for any finite set of formulas X. Note that there are only finitely many X-types realized in the universe of all models.

Definition 7.1 *An X-type is a finite set Σ of formulas from X satisfying the following properties for all formulas in the set X:*

$\neg\varphi \in \Sigma$ *iff* $\varphi \notin \Sigma$,

$(\varphi \vee \psi) \in \Sigma$ *iff* $\varphi \in \Sigma$ *or* $\psi \in \Sigma$.

if $[u/x]\psi \in \Sigma$, *then* $\exists x\psi \in \Sigma$.

Definition 7.2 *A type model $\mathfrak{M}$ for X is a finite set of X-types Σ satisfying the following existential witness property:*

if $\exists x\psi \in \Sigma$, then there is a type Δ in $\mathfrak{M}$ with (i) $\psi \in \Delta$ and (ii) the sets Σ and Δ contain exactly the same formulas whose free variables do not include x (written in what follows as: $\Sigma =_x \Delta$).

It is easy to see that each first-order model induces a type model.

Since there are only finitely many type models, it is decidable whether a given first-order formula has a type model. However, unlike in the case of modal logic, it should not (and in fact, it does not) follow that first-order logic is decidable. The reason is that not all type models as defined above come from standard models.

In this light, the filtration method as usually presented does two things at the same time: (a) it provides a finite ‘certificate’ for satisfiability in standard models, (b) the certificate is itself a standard model – something not needed for decidability.

7.2 Representation theorems for type models

To see how special it is to have a certificate for satisfiability that is itself a model, consider first-order logic. When can the above type models be represented as standard models? This may be a hard problem. Since first-order logic is undecidable, it must be undecidable whether a given type model is representable as having been induced by a standard model.[33]

[33] An interesting problem is determining the precise computational complexity of this notion.

But could not representability hold for suitable decidable fragments of first-order logic? This is indeed the case. The ‘guarded fragment’ GF of the first-order language was introduced in [3], who prove the following result for type models adapted to GF.

Theorem 7.3 *Type models for the guarded fragment can always be represented as standard models.*

The analogy with filtration is one of the many ways in which GF behaves like a modal logic. Even so, guarded type models are not standard first-order models themselves: crucially, they lack objects and assignments, and these have to be produced in the representation argument. This difference illustrates the earlier point about a finite certificate being good enough.[34]

At this point, a natural question arises. Even for the whole language, since the property of having a type model is decidable, what is the decidable first-order logic of all type models? Could not this logic be analyzed by broadening the class of standard models? The answer is positive, in terms of a broader semantic class of ‘generalized assignment models’, [2, 6], which have only a special subset of admissible assignments to be used in evaluation of quantifiers, not necessarily the family of all functions from variables to objects. [35]

Theorem 7.4 *Every type model is representable as a generalized assignment model.*

Proof. Consider any type model as described above. Define a *path* π as a finite sequence of types from the model, separated by variables x in such a way that each immediate x-transition in the path is from some type Σ to some type Δ with $\Sigma =_x \Delta$. Here *Last*(π) is the last type on the path π. Then define objects as pairs (π, x) of a path π and a variable x such that the last change on the path was via $=_x$. Next, for predicates P, set $I(P)((\pi_1, x1), \ldots, (\pi_k, x_k))$ iff (a)$\pi_1, \ldots, \pi_k$ all lie on one sequence where one is the longest, (b) no variables for the mentioned objects changed their variables later on along this longest path, (c) the atom $Px_1 \ldots x_n$ is a member of the last type on the path. Finally, the associated assignment Ass(π) for a path assigns to each variable x the object (π', x), where π' is an initial subpath of π and x did not change its value at some stage on the path π. Taken together, this gives a set of objects, an interpretation function, and a special set of assignments that form a generalized assignment model $\mathfrak{M}$ for first-order logic. $\square$

Here is the crucial property of this construction.

Fact 7.5 $\mathfrak{M}$, *Ass*(π) $\models \varphi$ *iff* $\varphi \in$ *Last*(π).

Proof. The proof is by induction on the formulas φ, where the two Boolean steps are routine by the properties of types. The atomic step plus the inductive step for the existential quantifier explain our choice of objects in $\mathfrak{M}$.

---

[34] In fact, GF does have the FMP, [49, 50], but the argument establishing that is much more graph-theoretic and combinatorial.

[35] In such models, with a variable assignment s, $\mathfrak{M}, s \models \exists x \varphi$ iff there exists an object d in the domain of $\mathfrak{M}$ such that (i) the assignment $s[x := d]$ is admissible, and (ii) $\mathfrak{M}, s[x := d] \models \varphi$.

For a concrete illustration of the base case, consider an atom Rxy. Let the values of Ass(π) on x,y be the objects (π_1,x) and (π_2,y), where by the linearity condition on admissible assignments, π_1,π_2 are subpaths of π. If $\mathfrak{M}$, Ass(π) $\models Rxy$, then $I(R)((\pi_1,x),(\pi_2,y))$. By definition, this means that the formula Rxy is an element of the last type on the longest of π_1,π_2, and also, none of the variables x,y changed their value on π afterwards. It follows that $Rxy \in Last(\pi)$. The argument in the opposite direction is similar.

Next, consider the inductive step for an existential quantifier $\exists x\psi$.

Case (i). If $\exists x\psi \in$ Last(π), then by the closure condition for existential quantifiers in type models, there exists a type Σ which contains ψ and which agrees with Last(π) on all x-free formulas. Consider the path π^+ consisting of π with Σ appended. By the inductive hypothesis, it holds that $\mathfrak{M}$, Ass(π^+) $\models \psi$. Therefore, $\mathfrak{M}$, Ass(π^+) $\models \exists x\psi$. Moreover, given the condition on equality of formulas, no changes in the other free variables $\bar{y}$ have taken place from the last type of π^+ to the last change made to these free variables from where they were last changed in π. Therefore, Ass(π) and Ass(π^+) agree on the values for these variables, and so we have that $\mathfrak{M}$, Ass(π) $\models \exists x\psi$.

Case (ii). If $\mathfrak{M}$, Ass(π) $\models \exists x\psi$, then there is an object d such that (a) $\mathfrak{M}$, Ass(π)$[x := d] \models \psi$, and (b) the assignment Ass(π)$[x := d]$ is admissible in the model. The latter fact means that it is of the form Ass(π') for some path π', where d is an object of the form (π'',u) for some subsequence π'' of π', where the variable u does not change its value any more towards the end.[36] By the inductive hypothesis, we have that $[u/x]\psi \in$ Last(π'). It follows that also $[u/x]\psi \in$ Last(π) by carefully analyzing which variables do not change their values along the paths π' and π.[37] Finally, by the relevant closure condition on types, if $[u/x]\psi \in$ Last(π), then $\exists x\psi \in$ Last(π) – and we are done.[38] □

Corollary 7.6 *The complete decidable logic of type models is the logic of generalized assignment models, which is known to consist of the following principles:*

1. *the minimal modal logic part of first-order logic, [8],*
2. *S5 for each quantifier $\exists x$, plus*
3. *all implications $(\neg)P\bar{x} \to \forall\bar{y}(\neg)P\bar{x}$ with $\bar{x},\bar{y}$ disjoint sequences of variables.* [39]

This may be viewed as modalizing the whole semantics of the complete first-order language. Cf. [2, 11] for more on this perspective.

---

[36] This is the crucial point in the proof where we use the restriction to admissible assignments.

[37] In particular, none of the paths π and π' needs to be a continuation of the other: they can fork off at some shared initial subpath. However, the argument about variables not changing their values, and the matching preservation of formulas in these variables also works in this extended setting, by going down to the forking point and then back up.

[38] The representation argument would also work, with some obvious modifications, for polyadic quantifiers $\exists\bar{x}$ over tuples of objects. On general assignment models, these are not reducible to iterated single-variable quantifiers.

[39] This result can also be proved by a reduction to the Guarded Fragment in the style of [3].

Remark 7.7 Generalized assignment semantics also works for extended first-order languages with irreducible polyadic quantifiers and substitution operators. This may reflect the flexibility in extending basic modal filtration to richer modal languages.

7.3 Conclusion

Filtration as a method for establishing decidability finds its boundaries in first-order logic. But filtration as a semantic style of thinking has a broader reach, suggesting even a new take on first-order semantics.

8 Summary and further directions

Filtration is an old method in modal logic which is still very much alive, and which keeps suggesting new perspectives and open problems. In this paper, we have surveyed the state of the art in classical filtration methods for various modal logics in various modal languages, pointing out its semantic and proof-theoretic aspects. After that, we discussed three more recent and perhaps unusual perspectives, namely, modalizing the theory of filtration, placing filtration inside a broader family of model updates, and connecting it to modal logics of questions and issues. Finally, we explored some boundaries, showing what filtration still has to say about generalized semantics for a first-order language. In the process, we found interesting new perspectives on filtration, as well as a large number of new open problems.

We elaborate a bit on the latter aspect. The technical themes presented here have by no means been exhausted. For instance, it should be possible to describe the syntactic formula sets used in filtration in more abstract terms, perhaps in terms of well-founded orderings on formulas that admit inductive analysis of the filtrated models. We also did not deeply explore other abstract takes on filtration, such as the analysis in [17, 18] using locally finite reducts and partial subalgebras.

Finally, in the realm of modal fixed-point logics, filtration seems close to the powerful automata techniques employed in the modal μ-calculus and related systems [33, 65], see also Gabbay's early use of Rabin's Theorem in modal logic [40, 41, 43].[40] The modal μ-calculus and its fragments are interesting systems for analysing the boundaries of the filtration method. As we have seen, PDL, a fragment of the μ-calculus, admits filtration, and [28] shows that filtration is applicable to non-trivial fragments of the modal μ-calculus other than PDL.[41] However, to prove the Finite Model Property for the full language of the modal μ-calculus, as

[40] Given that, by the Janin-Walukiewicz Theorem, the μ-calculus is the bisimulation-invariant fragment of Monadic Second-Order Logic, there may also be connections here with Fine's early work on second-order modal propositional logic, [34].

[41] The existence of a filtration method can even be seen as a criterion for simplicity of a logic.

shown by Kozen [52], filtration-like methods need to be considered together with the mathematical theory of well-quasi-orderings, suggesting another border line to be understood. For a recent analysis, cf. [1] on well-quasi-orderings and tree-like properties of infinitary proof systems for modal fixed-point logics.

The results in this paper also suggest interesting issues about understanding the scope of modal logic. For instance, we noted that filtration should fail for an undecidable system like first-order logic. But as we saw, that is not the end of the story. Modal perspectives are resilient, and filtration re-emerges in new forms when we tame complexity by remodeling the semantics of first-order logic, cf. [2]. As another example, consider the complexity of the meta-theory of modal logic discussed in Section 4. Working on the full meta-universe of modal models is working on a standard model whose complete theory can be very complex, but there could be principled reasons for restricting updates to 'available ones', as in the 'protocol models' of [12], thereby lowering complexity again in the style of [2].

But filtration is also attractive for general reasons beyond mathematical theory. Conceptually, it suggests going in the direction of 'relevance': representing just what is needed in semantic models, clearly a theme of general importance. And since what is needed may shift, there is also a need for connection between different levels of representation, suggesting a general quotient dynamics on different representation levels for information, or for reality.

In the latter vein, a technical 'old school' topic like filtration may even have a message for Kit Fine's current work on partiality and 'exact fit' in semantic modeling. The partiality used by many philosophical logicians concerns vocabulary, all infinitely many formulas over that restricted atomic vocabulary are allowed. In contrast, filtration may be called *super-partial*: it restricts the atoms, but also the relevant complex assertions one can make about them, and in doing so, it strikes an interesting balance between syntax and semantics.[42]

Acknowledgment We thank Bahareh Afshareh, Alexandru Baltag, Sebastian Enqvist and Julia Ilin, as well as the referees and the editors of this volume, for their helpful information and constructive comments.

References

1. B. Afshari, G. Jäger, and G. Leigh. An infinitary treatment of full mu-calculus. In *Proceedings of WoLLIC 2019*, LNCS. Springer, 2019.
2. H. Andrèka, J. van Benthem, N. Bezhanishvili, and I. Nèmeti. Changing a semantics: opportunism or courage? In M. Manzano et al., editor, *Life and work of Leon Henkin*, pages 307–337. Birkhäuser, 2014.

[42] Super-partiality is not the only philosophical take-home point from our technical analysis. For instance, also, the quasi-models found in Section 7 are a 'point-free' abstraction out of standard first-order models that may fit better with certain views of the universe in current metaphysics.

3. H. Andrèka, J. van Benthem, and Ì Nèmeti. Modal logics and bounded fragments of predicate logic. *Journ. Phil. Logic*, 27(3):217–274, 1998.
4. A. Baltag, N. Bezhanishvili, J. Ilin, and A. Özgün. Quotient dynamics: The logic of abstraction. In *International Workshop on Logic, Rationality and Interaction*, pages 181–194. Springer, 2017.
5. F. Bellissima. Atoms in modal algebras. *Math. Log. Q.*, 30(19-24):303–312, 1984.
6. J. van Benthem. *Exploring logical dynamics*. CSLI Publications, Stanford, 1996.
7. J. van Benthem. *Dynamic odds & ends*. Institute for Logic, Language and Computation (ILLC), University of Amsterdam, 1998.
8. J. van Benthem. *Modal Logic for open minds*. CSLI Publications, Stanford, CA, USA, 2010.
9. J. van Benthem. *Logical dynamics of information and interaction*. Cambridge University Press, Cambridge, UK, 2011.
10. J. van Benthem. Tracking information. In K. Bimbò, editor, *Michael Dunn on information–based logics*, pages 363–389. Springer Science Publishers, Dordrecht, 2016.
11. J. van Benthem and D. Bonnay. Modal logic and invariance. *J. Appl. Non-Classical Logics*, 18(2-3):153–173, 2008.
12. J. van Benthem, J. Gerbrandy, T. Hoshi, and E. Pacuit. Merging frameworks for interaction. *J. Phil. Logic*, 38(5):491–526, 2009.
13. J. van Benthem, P. Girard, and O. Roy. Everything else being equal: a modal logic for ceteris paribus preferences. *Journ. Phil. Logic.*, 38(1):83–125, 2009.
14. J. van Benthem and F. Liu. Dynamic logic of preference upgrade. *J. Appl. Non-Classical Logics*, 17(2):157–182, 2007.
15. J. van Benthem and S. Minica. Toward a dynamic logic of questions. *Journ. Phil. Logic*, 41(4):633–669, 2010.
16. J. van Benthem, B. ten Cate, and J. Väänänen. Lindström theorems for fragments of first-order logic. *Logical Methods in Computer science*, 5(3):67–96, 2009.
17. G. Bezhanishvili and N. Bezhanishvili. An algebraic approach to filtrations for superintuitionistic logics. In *A tribute to Albert Visser*, volume 30 of *Tributes*, pages 47–56. Coll. Publ., [London], 2016.
18. G. Bezhanishvili, N. Bezhanishvili, and R. Iemhoff. Stable canonical rules. *J. Symb. Log.*, 81(1):284–315, 2016.
19. G. Bezhanishvili, N. Bezhanishvili, and J. Ilin. Stable modal logics. *Rew. Symb. Logic*, 11(3):436–469, 2018.
20. G. Bezhanishvili and W. Holliday. A semantic hierarchy for intuitionistic logic. *Indagationes Mathematicae*, 30(2):403–469, 2019. special issue on “L.E.J. Brouwer, fifty years later”.
21. N. Bezhanishvili. *Lattices of Intermediate and Cylindric Modal Logics*. PhD thesis, University of Amsterdam, 2006.
22. N. Bezhanishvili, S. Ghilardi, and M. Jibladze. Free modal algebras revisited: the step-by-step method. In G. Bezhanishvili, editor, *Leo Esakia on Duality in Modal and Intuitionistic Logics*, pages 43–62. Springer, 2014.
23. N. Bezhanishvili and A. Kurz. Free modal algebras: A coalgebraic perspective. In Till Mossakowski, Ugo Montanari, and Magne Haveraaen, editors, *Algebra and Coalgebra in Computer Science, Second International Conference, CALCO 2007, Bergen, Norway, August 20-24, 2007, Proceedings*, volume 4624 of *Lecture Notes in Computer Science*, pages 143–157. Springer, 2007.
24. N. Bezhanishvili and B. ten Cate. Transfer results for hybrid logic. part I: the case without satisfaction operators. *J. Log. Comput.*, 16(2):177–197, 2006.
25. P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*. Cambridge University Press, 2001.
26. G. Boolos. *The logic of provability*. Cambridge university press, 1995.
27. R. A. Bull. That all normal extensions of S4.3 have the finite model property. *Z. Math. Logik Grundlagen Math.*, 12:341–344, 1966.
28. B. ten Cate and G. Fontaine. An easy completeness proof for the modal μ-calculus on finite trees. In C.-H. Luke Ong, editor, *Foundations of Software Science and Computational Structures, 13th International Conference, FOSSACS 2010*, pages 161–175, 2010.

29. A. Chagrov and M. Zakharyaschev. *Modal Logic*. The Clarendon Press, New York, 1997.
30. B. Chellas. *Modal logic: an introduction*. Cambridge university press, 1980.
31. I. Ciardelli, J. Groenendijk, and F. Roelofsen. *Inquisitive Semantics*. Oxford University Press, 2019.
32. W. Conradie, W. Morton, and C. van Alten. An algebraic look at filtrations in modal logic. *Log. J. IGPL*, 21(5):788–811, 2013.
33. S. Enqvist, F. Seifan, and Y. Venema. Completeness for the modal μ-calculus: separating the combinatorics from the dynamics. *Theoret. Comput. Sci.*, 727:37–100, 2018.
34. K. Fine. Propositional quantifiers in modal logic. *Theoria*, 36(3):336–346, 1970.
35. K. Fine. The logics containing S4.3. *Z. Math. Logik Grundlagen Math.*, 17:371–376, 1971.
36. K. Fine. Normal forms in modal logic. *Notre Dame Journal of Formal Logic*, 16(2):229–237, 1975.
37. K. Fine. Logics containing K4. II. *J. Symbolic Logic*, 50(3):619–651, 1985.
38. M. Fitting. Modal proof theory. In P. Blackburn, J. van Benthem, and F. Wolter, editors, *Handbook of modal logic*, volume 3 of *Stud. Log. Pract. Reason.*, pages 85–138. Elsevier B. V., Amsterdam, 2007.
39. D. Gabbay. A general filtration method for modal logics. *J. Philosophical Logic*, 30:323–330, 1970.
40. D. Gabbay. On decidable, finitely axiomatizable, modal and tense logics without the finite model property part i. *Israel Journal of Mathematics*, 10(4):478–495, 1971.
41. D. Gabbay. On decidable, finitely axiomatizable, modal and tense logics without the finite model property part ii. *Israel Journal of Mathematics*, 10(4):496–503, 1971.
42. D. Gabbay. Selective filtration in modal logic. *Theoria*, 1(1):29–34, 1972.
43. D. Gabbay. Decidability results in non-classical logics. *Annals of Mathematical Logic*, 8:237–295, 1975.
44. D. Gabbay. *Investigations into Modal and Tense logics, with Applications to Problems in Linguistics and Philosophy*. Reidel, Dordrecht, 1976.
45. S. Ghilardi. An algebraic theory of normal forms. *Ann. Pure Appl. Logic*, 71(3):189–245, 1995.
46. S. Ghilardi. Continuity, freeness, and filtrations. *Journal of Applied Non-Classical Logics*, 20(3):193–217, 2010.
47. J. Groenendijk and M. Stokhof. Questions. In J. van Benthem and A. ter Meulen, editors, *Handbook of Logic and Language*, pages 1055–1124. Elsevier/MIT Pess, 1997.
48. H. H. Hansen. *Monotonic modal logics*. Institute for Logic, Language and Computation (ILLC), University of Amsterdam, 2003.
49. I. Hodkinson. Loosely guarded fragment of first-order logic has the finite model property. *Studia Logica*, 70(2):205–240, 2002.
50. I. Hodkinson and M. Otto. Finite conformal hypergraph covers and Gaifman cliques in finite structures. *Bulletin of Symbolic Logic*, 9(3):387–405, 2003.
51. J. Ilin. *Filtration revisited: Lattices of stable non-classical logics*. PhD thesis, ILLC, University of Amsterdam, 2018.
52. D. Kozen. A finite model theorem for the propositional μ-calculus. *Studia Logica*, 47(3):233–241, 1988.
53. A. Kurz and A. Palmigiano. Epistemic updates on algebras. *arXiv preprint arXiv:1307.0417*, 2013.
54. E. J. Lemmon. Algebraic semantics for modal logics. I. *J. Symbolic Logic*, 31:46–65, 1966.
55. E. J. Lemmon. Algebraic semantics for modal logics. II. *J. Symbolic Logic*, 31:191–218, 1966.
56. E. J. Lemmon. *An Introduction to Modal Logic*. Basil Blackwell, Oxford, 1977. The "Lemmon notes", In collaboration with Dana Scott, Edited by Krister Segerberg, American Philosophical Quarterly, Monograph Series, No. 11.
57. I. Mason. The metatheory of the classical propositional calculus is not axiomatizable. *Journ. Symbolic Logic*, 50(2):451–457, 1985.
58. J. C. C. McKinsey. A solution of the decision problem for the Lewis systems S2 and S4, with an application to topology. *J. Symbolic Logic*, 6:117–134, 1941.

59. J. C. C. McKinsey and A. Tarski. The algebra of topology. *Annals of Mathematics*, 45:141–191, 1944.
60. L. Moss. Finite models constructed from canonical formulas. *J. Philosophical Logic*, 36(6):605–640, 2007.
61. E. Pacuit. *Neighborhood semantics for modal logic*. Springer, 2017.
62. K. Segerberg. *An Essay in Classical Modal Logic. Vols. 1, 2, 3*. Filosofiska Föreningen och Filosofiska Institutionen vid Uppsala Universitet, Uppsala, 1971.
63. O. Tuyt. *Canonical Rules on Neighbourhood Frames*. PhD thesis, Master's thesis, ILLC, Universiteit van Amsterdam, 2016.
64. A. Urquhart. Decidability and the finite model property. *J. Philosophical Logic*, 10(3):367–370, 1981.
65. I. Walukiewicz. Completeness of Kozen's axiomatisation of the propositional μ-calculus. *Inf. Comput.*, 157(1-2):142–182, 2000.
66. M. Zakharyaschev. Canonical formulas for K4. I. Basic results. *J. Symbolic Logic*, 57(4):1377–1402, 1992.
67. M. Zakharyaschev. Canonical formulas for K4. II. Cofinal subframe logics. *J. Symbolic Logic*, 61(2):421–449, 1996.