Dynamic vs. Classical Consequence

Denis Bonnay and Dag Westerståhl

Abstract The shift of interest in logic from just reasoning to all forms of information flow has considerably widened the scope of the discipline, as amply illustrated in Johan van Benthem's recent book *Logical Dynamics of Information and Interaction*. But how much does this change when it comes to the study of traditional logical notions such as logical consequence? We propose a systematic comparison between classical consequence, explicated in terms of truth preservation, and a dynamic notion of consequence, explicated in terms of information flow. After a brief overview of logical consequence relations and the distinctive features of classical consequence, we define classical and dynamic consequence over abstract information frames. We study the properties of information under which the two notions prove to be equivalent, both in the abstract setting of information frames and in the concrete setting of Public Announcement Logic. The main lesson is that dynamic consequence diverges from classical consequence when information is not persistent, which is in particular the case of epistemic information about what we do not yet know. We end by comparing our results with recent work by Rothschild and Yalcin on the conditions under which the dynamics of information updates can be classically represented. We show that classicality for consequence is strictly less demanding than classicality for updates.

Key words: Logical Consequence, Dynamic Semantics, Public Announcement Logic, Structural Rules, Logical Dynamics

---

Denis Bonnay
Département de Philosophie, Université Paris Ouest, 200 avenue de la République, 92000 Nanterre, France, e-mail: denis.bonnay@u-paris10.fr

Dag Westerståhl
Department of Philosophy, Stockholm University, SE-106 91 Stockholm, Sweden, e-mail: dag.westerstahl@philosophy.su.se

Johan van Benthem's recent book *Logical Dynamics of Information and Interaction* [7] can be seen as a passionate plea for a radically new view of logic. To be sure, the book is not a philosophical discussion of what logic is but rather an impressive series of illustrations of what logic *can be*, with presentations of numerous logical languages and a wealth of meta-logical results about them. The view is called simply Logical Dynamics, and contrasted with more traditional views of logic, and also with the earlier view from e.g. [4], now called Pluralism, in which logic was seen as the study of consequence relations.

According to Logical Dynamics, logic is not only about reasoning, about what follows from what, but about all aspects of *information flow among rational agents*. Not just proof and inference, but observations, questions, announcements, communication, plans, strategies, etc. are first-class citizens in the land of Logic. And not only the output of these activities belong to logic, but also the processes leading up to it.

This is a fascinating and inspiring view of logic. But how different is it from a more standard view? In particular, what does it change for the analysis of logical consequence, which had been the focus of traditional logical enquiry? This paper attempts some answers to the latter question, with a view to get clearer about the former.

1 Introduction

> ... in line with the thrust of this book, I see a discipline as a dynamic activity, not as any of its static products: proofs, formal systems, or languages. Logic is a stance, a modus operandi, and perhaps a way of life. That is wonderful enough. [7, p. 302]

Logicians will surely recognize this: doing logic is approaching your subject from a certain stance, a certain view of what the interesting questions are, what tools to use, what kind of abstractions are called for. The stance itself may be hard to put into words but is recognizable to the practitioners. Logical Dynamics is not really recommending a new stance towards logic, it seems to us. The novelty lies in what is taken to be its subject matter.

A rough object vs. meta level distinction is helpful: According to Logical Dynamics, many more object languages can and should be studied with logical methods than has traditionally been the case. But the form that this study takes is still of the familiar kind: (a) you introduce a formal language in which a particular variety of information flow can be expressed; (b) you provide a formal semantics and a deductive apparatus and see what can be proved from a choice of axioms; (c) you establish facts *about* these things: expressive power, definability, completeness, decidability, the structure of proofs, computational complexity, relations to other languages, etc. Indeed, this is precisely what a large part of *Logical Dynamics of Information and Interaction* is devoted to.

Reflection on logic from this perspective raises intriguing questions. What is it about a certain form of information flow that makes it apt for investigation by logical methods? What characterizes the syntactic constructs—the logical constants—used in the various object languages? Here we focus on just one aspect: the variety of consequence relations that emerge.

In fact, this variety is so great that it may in the end be just confusing to use the label "consequence relation" for all of them. Some have rather little to do with 'what follows from what' in any usual sense. But a distinction between two important kinds can be made: *dynamic* and *classical* relations. Very roughly: In order to draw the conclusion in the dynamic case, it may matter *how* the information from the premisses was processed. For example, during that process some information may get lost. For classical consequence, on the other hand, only the actual information contained in the premisses matters.

In particular, we are interested in how this distinction applies at the meta level. In contrast with the object level, where there seems to be no end to the variety of information-related activities that can be explored, the goals of meta level logical study seem rather fixed. One wants to know facts about the object level phenomena, that is, one wants the *truth* about them. Moreover, these truths are *mathematical* in a wide sense, and the only way you are allowed to assert a mathematical truth is to *prove* it. So the consequence relations operating at this level concern *reasoning towards truth*. This already separates them from a host of consequence relations resulting from object level phenomena. Does it in fact narrow the options down to just classical ones?[1]

To get a feeling for the issues involved here, and the variety of consequence relations on the market, let us look at a few examples.

- **Non-monotonicity 1**
  One way in which non-monotone consequence relations arise is when trying to model various kinds of reasoning under uncertainty, default reasoning, abductive reasoning, etc., where a conclusion is drawn tentatively, in awareness that it may have to be abandoned in the light of further evidence. There is no claim that the conclusion really *follows* from the premisses, and hence no real clash with classical consequence. Such reasoning is what everyone—even the logician—often has to resort to in daily life. But obviously it is never accepted in mathematics as conclusive grounds for a claim, and likewise not in metalogical reasoning. A number theorist may perhaps say that Goldbach's Conjecture is *likely* to be true, with reference to the so far observed even numbers greater than 2, but never that it is true (unless she has a proof).
- **Non-monotonicity 2**
  A different motive for rejecting monotonicity is proposed by *relevant* (or relevance) logicians. The idea is that adding an irrelevant premiss is not allowed. One question here concerns whether relevant logic is thought to be, in the terminology of John Burgess, *descriptive* (of the practice of mathematicians) or

[1] That is, classical in the sense just introduced. An intuitionistic consequence relation may well be classical in this sense.

prescriptive (wanting to reform that practice; see [8] and [20]). In any case, relevantists would seem to claim that relevance is or should be practiced at the meta level too.

But relevantists do aim at reasoning towards truth. And *no one* could argue that adding 'unnecessary' premisses, say in the form of a Weakening rule, threatens to lead from true premisses to false conclusions. To understand a relevantist stand on monotonicity, we need to be precise about which consequence relation is under discussion. Say that some relevant logic *Rel* is presented as a natural deduction system for deriving sequents of the form $\Gamma \succ \phi$. Then, if we define

$$\Gamma \vdash_{Rel} \phi \text{ iff there is a finite } \Gamma_0 \subseteq \Gamma \text{ s.t. } \Gamma_0 \succ \phi \text{ is derivable in } Rel,$$

monotonicity will hold.

Thus, for example, $\{p,q\} \vdash_{Rel} p$ (since $\{p\} \succ p$ and $\{p\}$ is a finite subset of $\{p,q\}$), even though $\{p,q\} \succ p$ may not be derivable. So a natural classical consequence relation extends the non-monotone relevantist one, in a way that can never be harmful, by anybody's lights, for reasoning towards truth.

- **Substructural logic**
  Weakening is a structural rule, and much recent work in logic rejects or modifies various such rules, e.g. in Linear Logic. But most of this work is not about truth at all. For a clear example, consider the Lambek Calculus (see e.g. [3]). Here Weakening, Contraction, and Permutation fail, for the obvious reason that the calculus aims to describe natural language syntax. Adding words, permuting words, or contracting two occurrences of the same word, may destroy well-formedness. This has nothing to do with truth, and indeed nicely illustrates the difference between consequence relations that logicians *study*, and the ones they may *use* in their own reasoning.
- **Contraction-free reasoning about truth**
  Consider a language for talking about truth: it contains a truth predicate, all instances of Tarski's T-schema, and some means for self-reference. With classical logic this leads to inconsistency. It has long been noted that proof-theoretic derivations of the Liar or Curry's Paradox rely on the *Contraction* rule,[2] and it has been suggested that dropping Contraction is a natural way to avoid paradoxes; for a recent proposal, see [22].
  This is a case where the difference between dynamic and classical consequence seems to matter. From a classical point of view, no one could seriously think that Contraction (as described in note 2) is an *invalid* rule. That would be like someone blaming your proof of a theorem B for using an assumption A *twice* (say, an earlier proved lemma), without explicitly saying so. First, you could retort that since A has been proved, you could easily repeat that proof twice in your proof of B. But really, the right answer is that the complaint makes no sense. When you claim that *B follows from A* (and possibly other assumptions), A and B are not *tokens*, although you need to use tokens of them when writing up the

[2] See e.g. [9]. We here intend a rule of the type "If $\Gamma, \phi, \phi \vdash \psi$ then $\Gamma, \phi \vdash \psi$". The validity of this rule need not entail the validity of, say, $\phi \rightarrow (\phi \rightarrow \psi) \vdash \phi \rightarrow \psi$.

proof. They are *types*, or *propositions*. And with these abstract objects it makes no sense to ask how many copies of them you are using.

But these considerations do not apply to a dynamic notion of consequence. An additional *order* is introduced: you process one thing *before* another thing, and then of course there is no guarantee that the usual structural rules will hold. Even if propositions are still types (not tokens), their processing happens in time, as it were.[3]

- **Public announcement**

  Our concern here, however, is not reasoning *about* truth, but reasoning with the aim of arriving at true propositions. Is there a role for dynamics here? Dynamic phenomena can themselves be expressed in richer logical languages. The clearest example is perhaps Public Announcement Logic (PAL) (introduced in [16]; see [7], ch. 3, for an overview). Here an announcement, when it can be made, changes the information state: situations where the announcement is false are discarded ('hard update'). Of course you cannot go around announcing anything: the claim has to be true in the current situation. What the announcement changes is agents' *knowledge*. In other words, this kind of reasoning only makes a difference when it is (also) about knowledge itself.

Although many non-classical consequence relations studied by logicians are irrelevant to our present concern with reasoning towards truth, the dynamic idea of consequence seems very different from the classical one. But note that rendering dynamic phenomena in a richer language like PAL is a *reduction* of dynamic consequence to classical consequence. That is, various dynamic consequence relations are expressible in PAL and similar logics, but PAL validity itself is classical. This is in fact part of the program in *Logical Dynamics of Information and Interaction*. The following quote is illustrative:

> Non-monotonicity is like a fever: it does not tell you which disease causes it. [7, p. 297]

Explicitly accounting for update phenomena in a richer language reveals the *causes* of non-monotonicity, but in a classical framework. The question is raised (but not answered) in the book whether this sort of reduction is always possible. That would be a very strong vindication of classical logic.

Our project in the rest of this paper is more modest. First, we specify precisely in what sense classical consequence is related to truth *preservation* (section 2). Then we compare classical and dynamic consequence in an update-friendly framework, both abstractly (section 3) and in the concrete setting of PAL (section 4), and show

[3] Zardini presents a formal system without Contraction containing a 'naive theory of truth' and shows that his consequence relation satisfies truth preservation in the following form (simplified): if $\Gamma, \phi \vdash \psi$, then $\Gamma \vdash Tr(\langle\phi\rangle) \rightarrow Tr(\langle\psi\rangle)$, where $\langle\phi\rangle$ names ϕ in the theory. By contrast, Field in [12] seems to agree with our point about Contraction and Permutation (ibid., pp. 10–11), but argues for an approach that rejects Excluded Middle (is 'paracomplete') as well as truth preservation in Zardini's sense. But note that this is a special version of truth preservation, tied to the occurrence of a truth predicate, and to the meaning of the conditional. As we will see in the next section, since Field's preferred consequence relation appears to be reflexive and transitive, there is a clear sense in which it necessarily preserves truth.

when the two notions of consequence coincide. Finally, we consider (section 5) a recent result in [17] on the conditions under which updates themselves are classical, and show that classicality of updates is a strictly stronger requirement than classicality of consequence.

2 Classical consequence and truth preservation

Following the remarks on Contraction and Permutation above, we take, in this section, consequence relations to hold between sets of sentences (the premisses) and sentences (the conclusion). The classical idea of logical consequence is as necessary (in some sense) truth preservation. There is no question that such reasoning is safe if your goal is arriving at the truth. But there are many different consequence relations that enforce truth preservation, for example, intuitionist as well as classical (in the sense of accepting excluded middle etc.) ones. So let us point out what they all have in common.

To begin, there may seem to be two ways to think about truth and truth preservation. One is in terms of *truth at a point* (think: possible world). The other is *truth in an interpretation*. But at the current abstract level, they are really equivalent.

To see this, fix a language L with its set $Sent_L$ of sentences. For truth at a point, suppose $\pi\colon Sent_L \to \wp(X)$ maps sentences to subsets of a set X of points, i.e. $\pi(\phi)$ is the set of points at which ϕ is *true*. Now consequence as necessary truth preservation (relative to π) is defined by

(1) $\Gamma \vdash^{\pi} \phi$ iff $\bigcap_{\psi\in\Gamma} \pi(\psi) \subseteq \pi(\phi)$

For the other approach to truth, let a *valuation* v assign truth values 1 or 0 to sentences, so we can take v to be a *subset* of $Sent_L$. Say that a sequent (Γ,ϕ) is *true in* v iff whenever $\Gamma \subseteq v$, we have $\phi \in v$. Note that sentence truth is a special case: ϕ is true in v iff $\phi \in v$ iff $(\emptyset,\phi)$ is true in v. For any set K of valuations we have a corresponding notion of truth-preserving consequence:

(2) $\Gamma \vdash_K \phi$ iff for all $v \in K$, (Γ,ϕ) is true in v.

To see that these approaches are equivalent, let π be given as above, and for each $a \in X$, let $v_a = \{\phi : \phi$ is true at $a\} = \{\phi : a \in \pi(\phi)\}$, and $K = \{v_a : a \in X\}$. Then:

(3) $\vdash^{\pi} = \vdash_K$

Conversely, if K is any set of valuations, let $\pi(\phi) = \{v \in K : \phi \in v\}$. Then again we obtain (3).

Since the two formats are equivalent, let us choose one: $\vdash_K$. Now stipulate that an arbitrary relation $\vdash$ between sets of sentences and sentences is *classical* iff it is reflexive and transitive in the following sense:

(R) If $\phi \in \Gamma$, then $\Gamma \vdash \phi$.

(T) If $\Delta \vdash \phi$, and for all $\psi \in \Delta$, $\Gamma \vdash \psi$, then $\Gamma \vdash \phi$.

Note that (R) and (T) entail monotonicity:[4]

(M) If $\Gamma \vdash \phi$ and $\Gamma \subseteq \Delta$, then $\Delta \vdash \phi$.

Now our point is this: reflexivity plus transitivity is *exactly* what constitutes necessary truth preservation in the above sense. This is a well-known fact, but let us spell it out.

Proposition 1. *$\vdash$ is classical if and only if $\vdash = \vdash_K$, for some K. In fact, we can take $K = Val(\vdash)$, so $\vdash$ is classical if and only if $\vdash = \vdash_{Val(\vdash)}$.*

Proof. It is obvious that each $\vdash_K$ satisfies (R) and (T). In the other direction, suppose $\vdash$ is classical. It is easy to verify from the definitions that $\vdash \subseteq \vdash_{Val(\vdash)}$. To prove the converse inclusion, suppose $\Gamma \not\vdash \phi$. Define the valuation v as follows: for any sentence ψ in L, $v(\psi) = 1$ iff $\Gamma \vdash \psi$. By (R), $v(\Gamma) = 1$, and by assumption, $v(\phi) = 0$. So it is enough to show that $v \in Val(\vdash)$. Suppose $\Delta \vdash \psi$ and $v(\Delta) = 1$; we must show that $v(\psi) = 1$. But it follows from (T) and the definition of v that $\Gamma \vdash \psi$, and we are done.

To repeat, this is well-known. For example, in view of (3), Proposition 1 is a reformulation of a representation theorem in [7, p. 297].[5] But the result is relevant to our discussion. Any instance of, say, non-monotonicity guarantees. *with respect to any class of interpretations*, that you will deduce a false conclusion from true premisses. That is just the trivial direction of Proposition 1. The slightly less trivial direction says that if you have Reflexivity and Transitivity, there is always at least one class of interpretations with respect to which you can construe your consequence relation as necessary truth preservation.[6]

Finally, let us emphasize again that adhering to truth-preservational consequence relations is perfectly compatible with intuitionistic or (most forms of) relevant consequence. That choice depends on whether you regard particular rules, such as $\neg\neg\phi \vdash \phi$ or $\phi, \neg\phi \vdash \psi$, as valid.[7]

---

[4] Another version of transitivity is Cut, in one of these two formulations:

(C1) If $\Delta \vdash \psi$ and $\Gamma, \psi \vdash \phi$, then $\Delta, \Gamma \vdash \phi$.
(C2) If $\Gamma \vdash \psi$ and $\Gamma, \psi \vdash \phi$, then $\Gamma \vdash \phi$.

If all sets are finite, we have (cf. [19], p. 17–18): (R)+(T) $\Leftrightarrow$ (C1)+(R)+(M) $\Leftrightarrow$ (C2)+(R)+(M).

[5] Also in [3], p. 247, and Prop. 7.4 in [4]. We are not sure who first made observation contained in Proposition 1. It appears in [18], presented without proof as a familiar fact, but apparently it goes back at least to [15]. The above proof (as well as a generalization of the result to multiple-conclusion logics) is given in [19] and in [14]. Indeed, Proposition 1 is a cornerstone in the abstract theory of consequence relations and propositional connectives expounded and elaborated in [14], especially chs. 1.1 and 3.

[6] We may note that the thesis of *Logical Pluralism* in [2] is in fact that logic studies classical consequence relations defined as in (1) (they call the points *cases*). In particular, they observe (though with some hesitation) that relevant consequence relations should be monotone.

[7] That, for example, intuitionistic propositional logic consequence $\vdash^{\mathrm{IL}}$ is classical in our sense means that Proposition 1 holds for it, i.e. $\vdash^{\mathrm{IL}}$ is determined by $Val(\vdash^{\mathrm{IL}})$. Since $\vdash^{\mathrm{IL}} \subseteq \vdash^{\mathrm{CL}}$, $Val(\vdash^{\mathrm{IL}})$

3 Two views about consequence

Judging from the previous section, analyzing consequence in terms of truth preservation for propositions does not seem to leave many options open. However, taking the intuitions of the dynamic perspective seriously, the picture becomes more complex. When asking whether something follows from some piece of information, one may ask whether it follows from information already secured and fully available, or whether it follows from information received along the way, which may not have been preserved. The first question corresponds to the classical notion of consequence: we look at where we are but not how we got there. In other words, we are not interested in how information has been received, but only in what follows from the information that we have. The second question corresponds to a thoroughly dynamic approach to consequence: rather than looking at where we are, we look at how we got there. In other words, we are interested in the information we received, which may end up not being what we have kept. Indeed, these two notions are likely to come apart because even hard information may not be preserved. Upon learning something which I did not know, I cease not to know it; I did receive new information, which allowed me to make some new inferences (e.g. that this something is true), but not all that information is being preserved.

We will now try to capture these two intuitions by means of formal definitions in an abstract dynamic setting.[8] Given a set of formulas L, an *abstract frame* for L is a structure

$$\mathscr{F} = (\Sigma, \{[\phi]\}_{\phi \in L})$$

where Σ is a set and each $[\phi] : \Sigma \to \Sigma$ is a partial function. Elements of Σ are to be construed as information states and each $[\phi]$ represents the effect of updating a given information state with the information that ϕ. Functionality expresses the fact that information updates are deterministic: the future state of information is completely determined by the past information state and the extra information that has been received. Partiality expresses the fact that not every piece of information is compatible with every information state: the (true) information that $\neg p$ is not compatible with the (truthful) announcement that p. Let the language L be fixed, together with a class of abstract frames $\mathfrak{F}$ representing the relevant possible informational scenar-

---

extends $Val(\vdash^{\mathrm{CL}})$ by allowing valuations that are not Boolean. More precisely, if $\mathscr{M}$ is any Kripke model for IL and $w \in |\mathscr{M}|$, the corresponding valuation $v_w^{\mathscr{M}}(\phi)$, consisting of the true sentences in $\mathscr{M}, w$, is equal to $Val(\vdash^{\mathrm{IL}})$ (by completeness), and $v_w^{\mathscr{M}}$ is Boolean for $\wedge$ and $\vee$, but not necessarily for $\neg$ or $\to$, since one may have e.g. $\phi \notin v_w^{\mathscr{M}}$ and $\neg\phi \notin v_w^{\mathscr{M}}$.

[8] Technically speaking, in such a setting, there would be more possible definitions of dynamic consequence than the two we are going to discuss—see [4] for more on this abstract stage setting, and an investigation into some more possibilities. But we take these two to be representative of the alternative between an essentially classical approach to consequence and an essentially dynamic one.

ios.[9] The thoroughly dynamic notion of consequence may be captured by what is known as Update to Test consequence:

Definition 1 (Update to Test consequence).
$\phi_1,\ldots,\phi_n \vDash^{\mathfrak{F}}_{UT} \psi$ iff for every $\mathscr{F} \in \mathfrak{F}$, $range([\phi_n] \circ \ldots \circ [\phi_1]) \subseteq fix([\psi])$.

(Here *range* assigns to every function in $\mathscr{F}$ its range and *fix* its set of fixed points.) This captures the dynamic notion because we ask whether we are in a fixed point for $[\psi]$ after shifting our information state along $[\phi_1],\ldots,[\phi_n]$. By contrast, the classical notion will be stated by considering an information state in which $[\phi_1],\ldots,[\phi_n]$ steadily hold; this is known in the dynamic literature as Test to Test consequence:

Definition 2 (Test to Test consequence). $\phi_1,\ldots,\phi_n \vDash^{\mathfrak{F}}_{TT} \psi$ iff for every $\mathscr{F} \in \mathfrak{F}$, $fix([\phi_1]) \cap \ldots \cap fix([\phi_n])\} \subseteq fix([\psi])$.

Indeed, this notion of consequence is classical in the sense of section 2: it satisfies (R) and (T).

Comparing these two notions, $\vDash^{\mathfrak{F}}_{UT} \subseteq \vDash^{\mathfrak{F}}_{TT}$ but the converse inclusion does not hold in general. As suggested earlier, it is bound to fail whenever information is not preserved. So when do classical and dynamic consequence come together? The intuitive answer is that they do so when the information represented behaves classically: received information is persistent, it gets into the current informational state and stays there.

To express this formally, a few definitions are needed. An abstract frame $\mathscr{F} = (\Sigma, \{[\phi]\}_{\phi\in L})$ is *idempotent* iff for all $\sigma \in \Sigma$, for all $\phi \in L$, $\sigma[\phi] = \sigma[\phi][\phi]$.[10] It is *commutative* iff for all $\sigma \in \Sigma$ and all $\phi,\psi \in L$, $\sigma[\phi][\psi] = \sigma[\psi][\phi]$. It is *f-commutative* iff for all $\sigma \in \Sigma$ and all $\phi,\psi \in L$, if $\sigma[\phi] = \sigma$, then $\sigma[\phi][\psi] = \sigma[\psi][\phi]$. So *f*-commutativity is a restricted version of commutativity, which only works for *f*ixed points and in the *f*orward direction.[11]

Idempotence and f-commutativity are conveniently made into a package deal:

Proposition 2. *An abstract frame $\mathscr{F} = (\Sigma, \{[\phi]\}_{\phi\in L})$ is idempotent and f-commutative iff for all $\phi \in L$, for any (possibly empty) sequence $\phi_1,\ldots,\phi_n$ of formulas in L,*

$$[\phi][\phi_1]\ldots[\phi_n] = [\phi][\phi_1]\ldots[\phi_n][\phi]$$

[9] Our definitions are relativized to a class $\mathfrak{F}$. Alternatively, we could have defined absolute notions by considering the greatest frame encompassing all possible informational scenarios. The relativized notions will help us make precise the role played by some technical assumptions.

[10] Here and below we use the arrow-like notation $\sigma[\phi] = \sigma'$, rather than the functional $[\phi](\sigma) = \sigma'$, to indicate that σ' is the result of updating σ with (the information that) ϕ:

$$\sigma \xrightarrow{\phi} \sigma'$$

Note that $[\phi][\psi]$ is the same as $[\psi] \circ [\phi]$. Note also that $[\phi]$ may be undefined for some σ. Throughout the paper, we take equalities $\sigma[\phi] = \sigma'[\psi]$ to mean that $\sigma[\phi]$ is defined iff $\sigma'[\psi]$ is, and that, when they are defined, they are equal.

[11] As pointed out to us by Seth Yalcin, *f*-commutativity can also be understood as a property of persistence of truth at a state.

Proof. From Left to Right: The proof is by induction on the length of the sequence of formulas. For the empty sequence, this is idempotence. Consider a sequence $\phi_1,\ldots,\phi_n,\phi_{n+1}$ and let $\sigma\in\Sigma$ be such that $\sigma[\phi][\phi_1]...[\phi_n][\phi_{n+1}]$ exists. By induction hypothesis, $[\phi][\phi_1]...[\phi_n]=[\phi][\phi_1]...[\phi_n][\phi]$. Hence we may apply f-commutativity at $\sigma[\phi][\phi_1]...[\phi_n]=\sigma'$ for ϕ and ϕ_{n+1}:

$$\sigma'[\phi][\phi_{n+1}]=\sigma'[\phi_{n+1}][\phi]$$

Replacing $\sigma'[\phi]$ by σ' in the left-hand side yields

$$\sigma[\phi][\phi_1]...[\phi_n][\phi_{n+1}]=\sigma[\phi][\phi_1]...[\phi_n][\phi_{n+1}][\phi]$$

as required.

From Right to Left: Idempotence is the case when the sequence is empty. As to f-commutativity, let $\sigma\in\Sigma$ be such that $\sigma[\phi]=\sigma$ and suppose $\sigma[\phi][\psi]$ exists. (Note that if we instead suppose $\sigma[\psi][\phi]$ exists, so does $\sigma[\psi]$, and hence, since $\sigma[\phi]=\sigma$, $\sigma[\phi][\psi]$ exists.) By hypothesis, $\sigma[\phi][\psi]=\sigma[\phi][\psi][\phi]$. Since $\sigma[\phi]=\sigma$, $\sigma[\phi][\psi][\phi]=\sigma[\psi][\phi]$, hence $\sigma[\phi][\psi]=\sigma[\psi][\phi]$ as required. □

The equality

$$[\phi][\phi_1]...[\phi_n]=[\phi][\phi_1]...[\phi_n][\phi]$$

intuitively means that the information that ϕ is persistent, in the sense that, once received, it still holds after updating in turn with $\phi_1,...,\phi_n$. Proposition 2 then says that, taken together, idempotence and f-commutativity precisely amount to information always being persistent. It will come as no surprise that our two notions of consequence coincide on the class of idempotent and f-commutative abstract frames:[12]

Proposition 3. *An abstract frame $\mathscr{F}$ is such that $\models_{UT}^{\{\mathscr{F}\}}=\models_{TT}^{\{\mathscr{F}\}}$ iff $\mathscr{F}$ is idempotent and f-commutative.*

Proof. If a frame $\mathscr{F}$ is such that $\models_{UT}=\models_{TT}$ (suppressing $\mathscr{F}$ in the notation), then it is idempotent and f-commutative: Since $\models_{TT}$ satisfies (R), so does $\models_{UT}$. Take $\phi,\phi_1,\ldots,\phi_n\in L$, $\sigma\in\Sigma$ and consider $\sigma[\phi][\phi_1]...[\phi_n]$, assuming it is defined. By (R), $\phi,\phi_1,...,\phi_n\models_{UT}\phi$. By definition of $\models_{UT}$, this means that $range([\phi_n]\circ...\circ[\phi_1]\circ[\phi])\subseteq fix([\phi])$. Hence $\sigma[\phi][\phi_1]...[\phi_n]=\sigma[\phi][\phi_1]...[\phi_n][\phi]$, and the result follows from Proposition 2.

If a frame is idempotent and f-commutative, then $\models_{UT}=\models_{TT}$: First, it is always the case that $\models_{UT}\subseteq\models_{TT}$, so all we need to prove is $\models_{UT}\supseteq\models_{TT}$. Assume $\phi_1,...,\phi_n\models_{TT}\psi$, and let $\sigma'\in range([\phi_n]\circ...\circ[\phi_1]$. There is $\sigma\in\Sigma$ such that $\sigma[\phi_1]...[\phi_n]=\sigma'$. Using idempotence and f-commutativity we see that $\sigma[\phi_1]...[\phi_n][\phi_i]=\sigma[\phi_1]...[\phi_n]$, i.e. $\sigma[\phi_1]...[\phi_n]\in fix([\phi_i])$, for every $i\in\{1,...,n\}$. Since $\phi_1,...,\phi_n\models_{TT}\psi$, this implies that $\sigma[\phi_1]...[\phi_n]\in fix([\psi])$. Thus $\phi_1,...,\phi_n\models_{UT}\psi$ as required. □

---

12 This result generalizes Proposition 2.3 in [21] by providing necessary as well as sufficient conditions for the equivalence to hold.

4 Classic and dynamic consequence in PAL

Proposition 3 provides an *abstract* characterization of the properties of information that make classical and dynamic consequence coincide. This abstract perspective can be made *concrete* by looking at specific forms of information update expressed in specific logical languages. This is the representation step advocated by [7]. We take such a step in the present section, and instantiate the two consequence relations within the logic of public announcements, PAL, where information updates are truthful public announcements.[13]

Technically, PAL expresses updates by means of an announcement operator $[!\_]$, which goes together with epistemic modalities K_i and a common knowledge modality C. Sets of models for PAL provide us with concrete versions of the abstract frames we were considering. More precisely, a *concrete frame* $\mathscr{K}$ is a set of multi-modal pointed S5-models $\mathscr{M},w$, which is closed under submodels and change of designated world. (Putting the two conditions together, if $\mathscr{M},w \in \mathscr{K}$ and $\mathscr{M}' \subseteq \mathscr{M}$, then $\mathscr{M}',w' \in \mathscr{K}$, for $w \in |\mathscr{M}|$ and $w' \in |\mathscr{M}'|$.) Closure under submodels guarantees that updates can be performed. The necessity to allow for changes in the designated world will become clear later.

Any concrete frame $\mathscr{K}$ generates an abstract frame

$$\mathscr{F}^{\mathscr{K}} = (\mathscr{K}, [\phi]\}_{\phi \in L})$$

where $[\phi]$ records the effect of updating with ϕ. Thus, $[\phi](\mathscr{M},w)$ is $\mathscr{M}|\phi,w$ if $\mathscr{M},w \vDash \phi$ and undefined otherwise, where $\mathscr{M}|\phi$ is $\mathscr{M}$ restricted to the worlds in $|\mathscr{M}|$ in which ϕ is true. Given a class of concrete frames $\mathfrak{K}$, we write $\mathfrak{F}^{\mathfrak{K}}$ for the class of abstract frames generated by frames in $\mathfrak{K}$. A PAL formula ϕ is valid on a concrete frame $\mathscr{K}$ iff it is true in every pointed model in $\mathscr{K}$, and it is valid on a class $\mathfrak{K}$ of concrete frames (notation: $\vDash^{\mathfrak{K}}$) iff it is valid on every frame in $\mathfrak{K}$. Also, observe that, for each $\mathscr{M},w \in \mathscr{K}$,

$$\begin{aligned}\mathscr{M},w \in \mathit{fix}([\psi]) &\Leftrightarrow \text{for all } w' \in |\mathscr{M}|,\ \mathscr{M},w' \vDash \psi \\ &\Leftrightarrow \mathscr{M} \vDash \psi\end{aligned}$$

We can now see the interplay between consequence relations and their concrete representations through the following two equivalences. First, as the label has it, Update to Test consequence is the abstract version of testing after updating:

Proposition 4. $\vDash^{\mathfrak{K}} [!\phi_1]\dots[!\phi_n]C\psi$ *iff* $\phi_1,\dots,\phi_n \vDash^{\mathfrak{F}^{\mathfrak{K}}}_{UT} \psi$.

This correspondence is known,[14] but we give a detailed proof here so as to make fully explicit what the assumptions are.

[13] Not all information updates are of this kind, e.g. because what we often get is 'soft' information that might be overridden. We leave a systematic investigation of the behavior of classical and dynamic consequence in these wider contexts to future research.

[14] [6], states that "modulo a few details, dynamic validity amounts to PAL validity" (p. 192). We spell out these details here.

Proof. From Left to Right: Assume that $\models^{\mathfrak{K}} [!\phi_1]\ldots[!\phi_n]C\psi$. Let $\mathscr{F}^{\mathscr{K}} \in \mathfrak{F}^{\mathfrak{K}}$ and $\mathscr{M},w \in \mathscr{K}$ be such that $\mathscr{M},w \in range([\phi_n] \circ \ldots \circ [\phi_1])$. We need to show that $\mathscr{M},w \in fix([\psi])$, that is, $\mathscr{M} \models \psi$. Let $w' \in |\mathscr{M}|$. Since $\mathscr{K}$ is closed under change of designated world, $\mathscr{M},w' \in \mathscr{K}$. Moreover, because the effect of an update does not depend upon which world is designated, $\mathscr{M},w' \in range([\phi_n] \circ \ldots \circ [\phi_1])$. By hypothesis, this guarantees that $\mathscr{M},w' \models C\psi$, hence in particular that $\mathscr{M},w' \models \psi$, as required.
From Right to Left: Assume that $\phi_1,\ldots,\phi_n \models^{\mathfrak{F}^{\mathfrak{K}}}_{UT} \psi$. Let $\mathscr{K} \in \mathfrak{K}$ and $\mathscr{M},w \in \mathscr{K}$; we need to show that $\mathscr{M},w \models [!\phi_1]\ldots[!\phi_n]C\psi$. That is, if the updates can be performed, ψ is common knowledge once they have been performed. So let $\mathscr{M}' = (\ldots(\mathscr{M}|\phi_1)\ldots)|\phi_n$. Thus, $\mathscr{M}',w \in \mathscr{K}$, and by definition of $\mathfrak{F}^{\mathfrak{K}}$, $\mathscr{M}',w \in range([\phi_n] \circ \ldots \circ [\phi_1])$. Therefore, by our initial assumption, $\mathscr{M}',w \in fix([\psi])$, so $\mathscr{M}' \models \psi$, which implies in turn that $\mathscr{M}',w \models C\psi$. □

Note that the left to right direction of the proof requires that frames are closed under change of designated world. If that were not the case, there could be a $\neg\psi$ world in the model which is not reachable from the designated world so that $C\psi$ holds even though the model is not a fixed point for updating with ψ. Since that part of the proof only uses the facticity of C, this also shows that $\models^{\mathfrak{K}} [!\phi_1]\ldots[!\phi_n]C\psi$ implies $\models^{\mathfrak{K}} [!\phi_1]\ldots[!\phi_n]\psi$, which is a valid rule in PAL.

Second, Test to Test consequence is the abstract version of a classical notion of consequence:

Proposition 5. $\models^{\mathfrak{K}} (C\phi_1 \wedge \ldots \wedge C\phi_n) \to C\psi$ *iff* $\phi_1,\ldots,\phi_n \models^{\mathfrak{F}^{\mathfrak{K}}}_{TT} \psi$.

Proof. The left to right direction is similar to the previous proof. Assume that $\models^{\mathfrak{K}} (C\phi_1 \wedge \ldots \wedge C\phi_n) \to C\psi$. Let $\mathscr{F}^{\mathscr{K}} \in \mathfrak{F}^{\mathfrak{K}}$ and $\mathscr{M},w \in \mathscr{K}$ be such that $\mathscr{M},w \in fix([\phi_i])$, $1 \leq i \leq n$. We need to show that $\mathscr{M},w \in fix([\psi])$, i.e. $\mathscr{M} \models \psi$. Let $w' \in |\mathscr{M}|$. Since $\mathscr{K}$ is closed under change of designated world, $\mathscr{M},w' \in \mathscr{K}$. And since $\mathscr{M} \models \phi_i$, we have $\mathscr{M},w' \models C\phi_i$ for $1 \leq i \leq n$. Sy by hypothesis, $\mathscr{M},w' \models C\psi$, from which it follows that $\mathscr{M},w' \models \psi$.
Right to Left: Assume $\phi_1,\ldots,\phi_n \models^{\mathfrak{F}^{\mathfrak{K}}}_{TT} \psi$. Let $\mathscr{K} \in \mathfrak{K}$ and $\mathscr{M},w \in \mathscr{K}$ be such that $\mathscr{M},w \models C\phi_1 \wedge \ldots \wedge C\phi_n$. We need to show that $\mathscr{M},w \models C\psi$. Let $\mathscr{M}^*$ be the submodel of $\mathscr{M}$ consisting of all those worlds that are connected to w. It is sufficient to show that $\mathscr{M}^*,w \models C\psi$. (Recall that the accessibility relations are equivalence relations.) Since $\mathscr{K}$ is closed under submodels, $\mathscr{M}^*,w \in \mathscr{K}$. Because $\mathscr{M}^* \models C\phi_1 \wedge \ldots \wedge C\phi_n$ and all worlds in $\mathscr{M}^*$ are connected to w, $\mathscr{M}^*,w \in fix([\phi_i])$, $1 \leq i \leq n$. Hence, by our initial assumption, $\mathscr{M}^*,w \in fix([\psi])$, which means that ψ is true everywhere in $\mathscr{M}^*$, so $\mathscr{M}^*,w \models C\psi$. □

Let *Prop* be the set of purely propositional formulas. It is well known that for such formulas classical consequence and its dynamic version coincide.

Proposition 6. *If* $\phi_1,\ldots,\phi_n,\psi \in Prop$, *then* $\models^{\mathfrak{K}} [!\phi_1]\ldots[!\phi_n]C\psi$ *iff* $\models^{\mathfrak{K}} (C\phi_1 \wedge \ldots \wedge C\phi_n) \to C\psi$.

Viewed from our perspective, Proposition 6 holds because of the special properties of propositional formulas with respect to updates. Propositional formulas generate idempotent and f-commutative frames. In the light of Proposition 3, together with Propositions 4 and 5, this readily implies Proposition 6. But this also suggests a more general question. Proposition 6 describes sufficient conditions for the concrete version of Proposition 3, which gives sufficient *and necessary* conditions. So, more generally, for which sublanguages of full modal logic do we have for every $\mathfrak{K}$ that $\vDash^{\mathfrak{K}} [!\phi_1]\ldots[!\phi_n]C\psi$ iff $\vDash^{\mathfrak{K}} (C\phi_1 \wedge \ldots \wedge C\phi_n) \rightarrow C\psi$? Proposition 3 says that we need to characterize the class of modal formulas that generate idempotent and f-commutative frames.

We shall consider this question for formulas of modal logic without common knowledge. This covers PAL without common knowledge by virtue of the reduction axioms. Let us say that a formula ϕ is *persistent* iff, for all concrete frames $\mathscr{K}$, $\mathscr{F}^{\mathscr{K}}$ restricted to $Prop \cup \{\phi\}$ (i.e. $(\mathscr{K}, \{[\psi]\}_{\psi \in Prop \cup \{\phi\}})$) is idempotent and f-commutative. On the face of it, generating idempotent and f-commutative frames is a property of sets of formulas. But giving a definition for formulas rather than sets thereof is adequate, because a set $\Gamma \supseteq Prop$ of formulas generates idempotent and f-commutative frames iff every formula in Γ is persistent in the sense just defined. (Testing persistence against *propositional* formulas is sufficient because the effect of updating with a non-propositional formula can always be mimicked using a suitably interpreted propositional formula.) Our question about formulas for which classical and dynamic consequence coincide may then be thus phrased:

Question 1. What is the class of persistent modal formulas?

Persistence can be analyzed into two more familiar features, corresponding respectively to f-commutativity and idempotence. A formula ϕ is *globally preserved under submodels* iff for any $\mathscr{M}$ and $\mathscr{M}'$ with $\mathscr{M}' \subseteq \mathscr{M}$, if $\mathscr{M} \vDash \phi$, then $\mathscr{M}' \vDash \phi$. A formula ϕ is *successful* if for any $\mathscr{M}, w$, if $\mathscr{M}, w \vDash \phi$, then $\mathscr{M}|\phi, w \vDash \phi$. Global preservation under submodels is indeed equivalent to the fact that, for any ψ, if $\mathscr{M}, w \in \mathit{fix}([\phi])$ and $\mathscr{M}', w = [\psi](\mathscr{M}, w)$ then $\mathscr{M}', w \in \mathit{fix}([\phi])$. Successfulness is equivalent in turn to the fact that $\mathit{range}([\phi]) \subseteq \mathit{fix}([\phi])$. (By the reasoning used in the proof of Proposition 4, success at a world makes for success at every world.)

Thus, Question 1 actually asks for a characterization of the class of modal formulas that are both successful and globally preserved under submodels. By Corollary 6.4 in [11], a formula is globally preserved under submodels iff it is globally equivalent to a universal modal formula.[15] The long sought-after characterization of successful formulas was provided in [13] where it is shown that, in S5, all unsuccessful formulas are look-alikes of the infamous Moore formula $p \wedge \neg Kp$. Note that if we were talking about local, instead of global, preservation under submodels,[16]

[15] A formula ϕ is globally equivalent to a universal modal formula if there is a formula ψ constructed using only (negations of) atoms, conjunction, disjunction and K such that for all $\mathscr{M}$, $\mathscr{M} \vDash \phi$ iff $\mathscr{M} \vDash \psi$.

[16] ϕ is locally preserved under submodels iff $\mathscr{M}' \subseteq \mathscr{M}$, $w \in |\mathscr{M}'|$, and $\mathscr{M}, w \vDash \phi$ implies $\mathscr{M}', w \vDash \phi$.

we would be done. As shown in [1], a formula is locally preserved under submodels iff it is locally equivalent to a universal formula (we will take the liberty to say 'locally universal'), and universal formulas are always successful. However, matters are more complicated when global preservation is concerned: in S5, the Moore formula is trivially globally preserved under submodels, since it is not globally satisfiable, but it is not successful.

Limiting ourselves from now on to a modal language with only one epistemic modality (as we in effect did in the previous paragraph), we provide a partial answer to Question 1. In order to do so, Carnapian disjunctive normal forms for modal formulas prove useful. (This strategy is inspired by [13].)

Definition 3. A formula δ is in *normal form* iff it is a disjunction of conjunctions of the form $\delta = \alpha \wedge \Box\beta_1 \wedge \ldots \wedge \Box\beta_n \wedge \Diamond\gamma_1 \wedge \ldots \wedge \Diamond\gamma_m$, where α and each γ_i are conjunctions of literals and each β_j is a disjunction of literals.

For the limited class of modal formulas such that their disjunctive normal form consists of only one disjunct, Question 1 gets a rather satisfactory answer.

Proposition 7. *Let δ be a conjunction in normal form. In S5, δ is persistent iff δ is locally universal.*

Proof. From Left to Right: We prove the contrapositive. Assume δ is not locally universal. Either δ is globally satisfiable (meaning that there is a $\mathscr{M}$ such that $\mathscr{M} \vDash \delta$) or it is not. If δ is not globally satisfiable, note first that it still is locally satisfiable (meaning that there is a $\mathscr{M}, w$ such that $\mathscr{M}, w \vDash \delta$), or δ would be locally equivalent to $p \wedge \neg p$, which is universal. Being locally but not globally satisfiable, it readily follows that δ is not successful, hence not persistent, and we are done. So we assume that δ is globally satisfiable. There has to be a γ_i such that $\alpha \wedge \Box\beta_1 \wedge \ldots \wedge \Box\beta_n \nvDash \gamma_i$, since otherwise δ would be equivalent to $\alpha \wedge \Box\beta_1 \wedge \ldots \wedge \Box\beta_n$, which is locally universal. It follows that

(*) $\quad \alpha \wedge \beta_1 \wedge \ldots \wedge \beta_n \wedge \neg\gamma_i$ is satisfiable.

Take a model $\mathscr{M}$ such that $\mathscr{M} \vDash \delta$. By (*), it is possible to extend $\mathscr{M}$ to a model $\mathscr{M} \cup \{w\}$ such that $\mathscr{M} \cup \{w\}, w \vDash \alpha \wedge \beta_1 \wedge \ldots \wedge \beta_n \wedge \neg\gamma_i$ and $\mathscr{M} \cup \{w\} \vDash \delta$. But then $\mathscr{M} \cup \{w\}|(\alpha \wedge \beta_1 \wedge \ldots \wedge \beta_n \wedge \neg\gamma_i) \nvDash \delta$, so δ is not globally preserved under submodels.

From Right to Left: It is known that locally universal formulas are successful, see e.g. [5]. Moreover, by the result in [1], universal formulas are locally preserved under submodels, and *a fortiori* globally so. $\Box$

Thus, for conjunctions in normal form, classical and dynamic notions of consequence are equivalent exactly when the information which comes into play is stable, in the sense of being locally preserved under submodels. We leave it as an open question whether this result carries over: is it the case that, in S5, for any formula ϕ, ϕ is persistent iff ϕ is locally universal?

5 Classical consequence *vs*. classical update

In the last two sections we studied the conditions under which abstract and concrete frames behave classically with respect to dynamic consequence. Rothschild and Yalcin in [17] recently asked similar questions, concerning the conditions under which abstract frames behave classically with respect to updates themselves. Comparing results will prove instructive. The main lesson of Proposition 7 is that some non-propositional formulas pass the classicality test for consequence. Could it be so for updates as well? The answer is not immediate. Intuitively, being classical with respect to updates is more demanding than being classical only with respect to the visible effects of these updates on the consequence relation.

Following Rothschild and Yalcin, being classical with respect to updates means that informational states and propositional contents can be represented by sets of worlds, in such a way that updating with ϕ is taking the intersection of the current informational state with the set of ϕ-worlds. When this is so, the abstract frame is said to be static:[17]

Definition 4 (Rothschild & Yalcin). An abstract frame $(\Sigma, \{[\phi]\}_{\phi \in L})$ is *static* iff there are functions $f : L \to \wp(\Sigma)$ and $g : \Sigma \to \wp(\Sigma)$ such that for any $\sigma \in \Sigma$, $g(\sigma[\phi]) = g(\sigma) \cap f(\phi)$.

A result similar to Proposition 3 ensues.

Proposition 8 (Rothschild & Yalcin). *An abstract frame is static iff it is idempotent and commutative.*

Clearly, commutativity implies f-commutativity, but the converse is not true, even assuming idempotence. Actually, idempotence and f-commutativity correspond to a weaker notion of being static where $g(\sigma[\phi]) = g(\sigma) \cap f(\phi)$ is replaced by $g(\sigma[\phi]) \subseteq g(\sigma) \cap f(\phi)$ in Definition 4.

Going concrete, we shall say that a modal formula ϕ is *strongly persistent* iff, for all concrete frames $\mathscr{K}$, $\mathscr{F}^{\mathscr{K}}$ restricted to $Prop \cup \{\phi\}$ is idempotent and commutative. (Just as with 'persistent', the definition of 'strongly persistent' can be given for formulas rather than sets of formulas, and for similar reasons.) Here is a full characterization of the strongly persistent formulas.

Proposition 9. *Let ϕ be any modal formula. In S5, ϕ is strongly persistent iff ϕ is equivalent to a propositional formula.*

Proof. The direction from right to left is immediate. We prove the other direction.[18] Consider an arbitrary strongly persistent formula ϕ. First, we note that it is sufficient to prove the following:

[17] This exact definition was to be found in an early version of [17]. The definition in the final version is slighly more complex, but the extra complexity is irrelevant to our present purpose.

[18] Our proof of Proposition 9 is inspired by a simplified proof of Proposition 8 by Johan van Benthem (private correspondence). Whether there is deeper connection still remains to be seen. Is there a sense in which the possibility of a classical representation forces the equivalence to purely propositional formulas?

(*) If $\mathscr{M}, w \models \phi$ and $\mathscr{M}, w \equiv_{Prop} \mathscr{M}', w'$, then $\mathscr{M}', w' \models \phi$,

where $\equiv_{Prop}$ is elementary equivalence restricted to propositional formulas. To see this, let ϕ^{Prop} be the set of propositional consequences of ϕ. We claim that

(**) $\phi^{Prop} \models \phi$

By compactness, ϕ is then equivalent to a propositional formula. To prove (**), assume $\mathscr{M}', w' \models \phi^{Prop}$ and show $\mathscr{M}', w' \models \phi$. Let $Prop_{w'}$ be the set of propositional formulas that are true in $\mathscr{M}', w'$. $Prop_{w'} \cup \{\phi\}$ is consistent, since otherwise there would be propositional formulas $\psi_1, \ldots, \psi_n$ in $Prop_{w'}$ such that $\phi \models \neg(\psi_1 \wedge \ldots \wedge \psi_n)$, i.e. $\phi \models \neg\psi_1 \vee \ldots \vee \neg\psi_n$, contradicting the fact that $\mathscr{M}', w' \models \phi^{Prop}$. So there is $\mathscr{M}, w$ with $\mathscr{M}, w \models Prop_{w'}$ and $\mathscr{M}, w \models \phi$. $\mathscr{M}, w \models Prop_{w'}$ means that $\mathscr{M}, w \equiv_{Prop} \mathscr{M}', w'$ so (*) applies and we have $\mathscr{M}', w' \models \phi$, as required.[19]

We now prove (*), assuming that ϕ is strongly persistent. Consider two structures $\mathscr{M}, w$ and $\mathscr{M}', w'$ in a concrete frame $\mathscr{K}$ with $\mathscr{M}, w \models \phi$ and such that $\mathscr{M}, w \equiv_{Prop} \mathscr{M}', w'$. Since they are S5-models, we may also assume without loss of generality that there are no two different worlds satisfying exactly the same propositional formulas in $\mathscr{M}$ or $\mathscr{M}'$.

Since $\mathscr{M}, w \models \phi$, $[\phi]$ is defined on $\mathscr{M}, w$ in $\mathscr{F}^{\mathscr{K}}$. Let $Prop_w$ be the conjunction of propositional atoms and negations thereof that are true in $\mathscr{M}, w$. (The proof does not go through if there is an infinite number of atoms.) $[Prop_w]$ is also defined on $(\mathscr{M}, w)[\phi]$. By commutativity, $(\mathscr{M}, w)[\phi][Prop_w] = (\mathscr{M}, w)[Prop_w][\phi]$. Hence, by idempotence,

(***) $(\mathscr{M}, w)[\phi][Prop_w] = (\mathscr{M}, w)[\phi][Prop_w][\phi]$

Since $\mathscr{M}, w \equiv_{Prop} \mathscr{M}', w'$, $[Prop_w]$ is defined on $\mathscr{M}', w'$ too. Moreover, we have $(\mathscr{M}', w')[Prop_w] = (\mathscr{M}, w)[\phi][Prop_w]$, since the result of a successful announcement of $Prop_w$ is always the same one-world structure. Also, it follows from (***) that $(\mathscr{M}', w')[Prop_w] = (\mathscr{M}', w')[Prop_w][\phi]$. But then, by commutativity, $(\mathscr{M}', w')[\phi][Prop_w]$ is defined, which implies that $\mathscr{M}', w' \models \phi$. This completes the proof. □

In the context of PAL, classicality *with respect to consequence only* and classicality *with respect to updates in general* end up being two very different things. In the second case, only propositional formulas pass the test. Requiring classicality with respect to updates in general really means dealing with announcements that are deprived of epistemic content. By contrast, in the first case, universal formulas do pass the test. Requiring classicality with respect to consequence is compatible with information that has epistemic content, as long as this epistemic content is about knowledge rather than doubts.

[19] The first part of the proof is well-known from model theory, relying only on compactness and the fact that ϕ^{Prop} is closed under disjunction; see [10], Lemma 3.2.1.

6 Conclusion

What does logical dynamics tell us about logical consequence? Our focus has been properties of content and how they should be represented. Such questions are familiar in formal semantics, regarding whether certain phenomena (anaphora, presuppositions, etc.) should be given a classical account in terms of propositions interpreted as sets of possible worlds or dynamically represented in terms of context change potentials. Debates regarding logical consequence typically take a different form. The focus is on the validity of inference rules, and the model-theoretic definition of logical consequence most often remains classical. But, in principle, there is no telling apart the question of content and the question of validity: some rules may be valid only with respect to some particular types of contents. From the dynamic perspective, this becomes clear with the splitting of logical consequence into a classical and a dynamic notion. The broadest notion is the dynamic one, which is about what follows from the information received. In this context, classicality for logical consequence emerges as a property of content: when information is persistent, dynamic consequence boils down to the standard semantic definition in terms of preservation of truth and satisfies the classical structural rules.

This new take on logical consequence is also a new take on the dynamic *versus* classical dispute in semantics. The dispute has several faces: dynamic or classical *what*? Asking the question about updates is not the same as asking the question about logical consequence. Rothschild and Yalcin conclude their paper [17] by asking about classes of semantic systems which would lie between the static systems and the information-sensitive systems. Our static systems (the idempotent and f-commutative frames) constitute such a class: they may be information-sensitive for updates, but they are information-insensitive for logical consequence. This suggests a fully parametric approach: given a certain kind of manifestation of content (through logical consequence, updates of common ground, but also possibly other manifestations, such as, say, presupposition accommodation), what are the properties of content that make a classical analysis possible or force a dynamic account?

Acknowledgements Thanks to Chris Barker, Johan van Benthem, Rosalie Iemhoff, Daniel Rothschild and Seth Yalcin for helpful comments and suggestions. Both authors were supported by the ESF-funded project Logic for Interaction (LINT; http://www.illc.uva.nl/lint/index.php), a Collaborative Research Project under the Eurocores program LogICCC, and by a grant from the Swedish Research Council.

References

1. Hajnal Andréka, Johan van Benthem, and István Németi. Modal languages and bounded fragments of predicate logic. *Journal of Philosophical Logic*, 27:217–274, 1998.
2. JC Beall and Greg Restall. *Logical Pluralism*. Clarendon Press, Oxford, 2006.
3. Johan van Benthem. *Language in Action. Categories, Lambdas, and Dynamic Logic*. North-Holland/Elsevier, Amsterdam, 1991.

4. Johan van Benthem. *Exploring Logical Dynamics*. CSLI Publications, Stanford, 1996.
5. Johan van Benthem. What one may come to know. *Analysis*, 64:95–105, 2004.
6. Johan van Benthem. Logical dynamics meets logical pluralism? *Australasian Journal of Logic*, 6:182–209, 2008.
7. Johan van Benthem. *Logical Dynamics of Information and Interaction*. Cambridge University Press, Cambridge, 2011.
8. John P. Burgess. No requirement of relevance. In S. Shapiro, editor, *The Oxford Handbook of Philosophy of Mathematics and Logic*, pages 727–751. Oxford University Press, Oxford, 2005.
9. Andrea Cantini. Paradoxes and contemporary logic. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Summer 2012 edition, 2012.
10. C. C. Chang and H. Jerome Keisler. *Model Theory*. North-Holland, Amsterdam, 1973. Studies in Logic and the Foundations of Mathematics, v. 73, 3rd ed. 1990.
11. Maarten de Rijke and Holger Sturm. Global definability in basic modal logic. In H. Wansing, editor, *Essays on Non-Classical Logic*, pages 111–135. King's College University Press, London, 2000.
12. Hartry Field. *Saving Truth from Paradox*. Oxford University Press, Oxford, 2008.
13. Wesley H. Holliday and Thomas F. Icard. Moorean phenomena in epistemic logic. In *Advances in Modal Logic*, pages 178–199, 2010.
14. Lloyd Humberstone. *The Connectives*. MIT Press, Cambridge, MA, 2011.
15. Jerzy Łoś and Roman Suszko. Remarks on sentential logics. *Indagationes Mathematicas*, 20:171–183, 1958.
16. Jan A. Plaza. Logics of public communication. In *Methodologies for Intelligent Systems, 4*, pages 201–216. North-Holland, Amsterdam, 1989.
17. Daniel Rothschild and Seth Yalcin. On the dynamics of conversation. *manuscript*, 2012.
18. Dana Scott. On engendering an illusion of understanding. *The Journal of Philosophy*, 68:21:787–807, 1971.
19. D. J. Shoesmith and T. J. Smiley. *Multiple-Conclusion Logic*. Cambridge University Press, Cambridge, 1978.
20. Neil Tennant. Relevance in reasoning. In S. Shapiro, editor, *The Oxford Handbook of Philosophy of Mathematics and Logic*, pages 696–725. Oxford University Press, Oxford, 2005.
21. Jaap van der Does, Willem Groeneveld, and Frank Veltman. An update on might. *Journal of Logic, Language and Information*, 6(4):361–380, 1997.
22. Elia Zardini. Truth without contra(di)ction. *The Review of Symbolic Logic*, 4:04:498–535, 2011.