

Belief Contraction in Dynamic Epistemic Logic

Gaia Belardinelli

Department of Philosophy
Stanford University
Stanford, USA
gaiabel@stanford.edu

Snow Zhang

Department of Philosophy
University of California, Berkeley
Berkeley, USA
snowzhang@berkeley.edu

Dynamic epistemic logic (DEL) represents belief change via model transformations induced by epistemic events. Its standard formulation (Baltag, Moss, Solecki, 1998) provides a natural account of belief expansion through the elimination of possibilities, but it cannot model belief contraction about factual propositions. A classic response enriches Kripke models with plausibility orderings, representing contraction as an update that promotes certain possibilities over others. We show that this approach has expressive limitations. In particular, the approach cannot model belief that violates positive introspection and contraction dynamics in response to a *hedged* public announcement that φ *might* be false. Motivated by these considerations, we introduce a mechanism for belief contraction defined directly on standard Kripke models, without any constraints on the doxastic accessibility relation. We show that it satisfies some of the standard properties of belief contraction but not others, study the conditions under which contraction may be unsuccessful, and provide a sound and complete axiomatization of the logic via reduction axioms. We also define a more general dynamic logic that is an extension of standard DEL and accommodates belief contractions due to events such as private or semi-private announcements, and provide a complete and sound axiomatization of the general logic.

1 Introduction

Dynamic epistemic logic (DEL) is a family of logics that model multi-agent belief change by incorporating information revealed by epistemic events—such as public or private announcements—into agents’ epistemic states. In standard DEL [6], this incorporation proceeds by *eliminating* all possibilities that are incompatible with the disclosed information. As a result, standard DEL can’t straightforwardly model announcements that make agents reconsider possibilities they previously ruled out as impossible, as required for belief contraction. A popular solution is to distinguish “hard” from “soft” updates; while hard updates involve *eliminating* possibilities, soft updates involve *promoting* possibilities in terms of their comparative *plausibility* [12, 8]. Then, belief contraction about a factual proposition p consists in coming to judge that a $\neg p$ -possibility is at least as plausible as the most plausible p -possibility, thereby losing the belief that p holds.

While the plausibility-based approach can model belief contraction [8, 23], it has certain expressive limitations. First, since plausibility orderings are transitive, the framework validates positive introspection for beliefs (axiom 4), and so can’t model agents who are mistaken about their own beliefs. Second, the approach can’t accommodate belief contractions due to *hedged public announcements*—announcements that p *might be false*. Intuitively, such an announcement could cause an agent who initially believes p to suspend judgment on p without gaining any new (factual) beliefs. We show that there is a precise sense in which such dynamics *can’t* be modeled in the plausibility framework.

These limitations motivate the search for a new model of belief contraction. We offer such a model in this paper. We introduce a dynamic operation of belief contraction on standard, multi-agent Kripke

models of beliefs, with no assumptions on the doxastic accessibility relations. Roughly, a hedged announcement that φ *might be false* triggers a contraction on φ , which involves adding $\neg\varphi$ -possibilities to an agent's belief set if she believes φ , and doing nothing otherwise. We define a sound and complete logic for this contraction modality, which we call *hedged public announcement logic* (HPAL), and prove that this update satisfies a number of natural properties of belief contraction. We then define a more general dynamic logic GDEL that conservatively extends standard DEL to also accommodate belief contractions due to events such as private or semi-private announcements. We show that it contains HPAL and that its updates are always a simulation of a refinement of an initial Kripke model. Finally, we provide a complete and sound axiomatization of the general logic.

Related literature. The logic GDEL can be viewed as a natural extension of standard DEL. As contraction is a form of information-loss, our model is also related to logics of *forgetting* and simulation modal logic [18, 19, 22], as well as to the fragment of graph modifier logic that adds edges to Kripke models [5]. Our theory however contrasts with the theory of belief contraction as studied in AGM [1]; our contraction operator does not satisfy many of the AGM postulates. This is to be expected, as dynamic contraction involves transformations that could change the truth-values of epistemic formulas, which is not considered in AGM. Philosophically, the problem of belief contraction due to hedged public announcement is related to the problem of evaluating indicative conditionals that have modal antecedents, which remains an open problem [32, 27, 10].

The proofs of main theorems can be found in the Appendix.

2 Standard DEL framework

We begin by reviewing the basics of DEL and fix notations.¹ Let At be a countable set of atomic formulas and Ag be a finite set of agents. Let $\mathcal{L}$ be the *doxastic language* generated from At by the following grammar: $\varphi := \top \mid p \mid \neg\varphi \mid \varphi \wedge \varphi \mid B_a\varphi$, where $p \in At$ and $a \in Ag$. We write $\perp$ for $\neg\top$ and $\tilde{B}_a\varphi$ for $\neg B_a\neg\varphi$. A *literal* is an atomic formula p or its negation $\neg p$. Let $\bigwedge_{\varphi \in S} \varphi = \top$ if S is empty.

We adopt the standard semantics of epistemic logic, where a Kripke model is a triple $\mathcal{M} = (W, R, V)$ with $W \neq \emptyset$ a set of worlds, $R_a \subseteq (W \times W)$ an accessibility relation defined for all $a \in Ag$, and $V : At \rightarrow \wp(W)$ is a valuation function. We use notation $R_a(w) = \{v \in W : (w, v) \in R_a\}$ to represent *the set of possible worlds compatible with agent a 's belief at w* .

The satisfaction relation between a Kripke model $\mathcal{M} = (W, R, V)$ with world $w \in W$ and a formula $\varphi \in \mathcal{L}$ is defined as standard, where in particular the semantics of the belief modality is given by:

$$\mathcal{M}, w \models B_a\varphi \quad \text{iff} \quad \mathcal{M}, v \models \varphi \text{ for all } v \in R_a(w).$$

A formula φ is *valid in a model* $\mathcal{M}$ (notation: $\mathcal{M} \models \varphi$), if for all worlds w in $\mathcal{M}$, $\mathcal{M}, w \models \varphi$. A formula φ is *valid in a class of models* $\mathfrak{L}$ (notation: $\models_{\mathfrak{L}} \varphi$), if for all models $\mathcal{M}$ in $\mathfrak{L}$, φ is valid in $\mathcal{M}$.

A *standard event model* is a triple $\mathcal{E} = (E, Q, pre)$ where $E \neq \emptyset$ is a finite set of events, $Q_a \subseteq (E \times E)$ is an accessibility relation between events, for all $a \in Ag$, and $pre : E \rightarrow \mathcal{L}$ a precondition function.² We use notation $Q_a(e) = \{f \in E : (e, f) \in Q_a\}$ for the set of events that are considered possible by a at e .

For a Kripke model $\mathcal{M} = (W, R, V)$ and a standard event model $\mathcal{E} = (E, Q, pre)$, the *standard product update* of $\mathcal{M}$ with $\mathcal{E}$ is the Kripke model $\mathcal{M} \otimes \mathcal{E} = (W^{\mathcal{E}}, R^{\mathcal{E}}, V^{\mathcal{E}})$ where $W^{\mathcal{E}} = \{(w, e) \in W \times E : w \models$

¹For presentation purposes, in this section we only introduce the basic doxastic language without dynamic modalities. We introduce the full language of DEL in the last section, where we also use it. For an introduction to DEL, see [7].

²In the most general form of DEL, the precondition function assigns to each event a formula from the dynamic language (i.e. $\mathcal{L}$ extended with dynamic formulas). We keep it simple here and consider this general DEL form in the last section.

$pre(e)\}$, $R_a^\mathcal{E}(w, e) = \{(v, f) \in W^\mathcal{E} : v \in R_a(w) \text{ and } f \in Q_a(e)\}$, and $V^\mathcal{E}(p) = \{(w, e) \in W^\mathcal{E} : w \in V(p)\}$, for all $p \in At$. Informally, a world (w, e) represents *the state of affairs at w after e occurs*. We call these event models and product updates *standard* to distinguish them from the framework introduced later.

It is well known that standard DEL has undesirable consequences when agents are presented with information that contradicts their beliefs [12, 26]. For a concrete illustration:

Example 2.1. Alice believes that she locked her office door before she left. She runs into Bob, who tells her that *her office door was open*. Alice trusts Bob and revises her beliefs accordingly.

Let p denote the proposition “Alice’s office door is open”. Figure 1 illustrates that, if we model Alice’s belief revision using a standard product update, then she will end up having inconsistent beliefs.



Figure 1: **Left:** The Kripke model $\mathcal{M} = (W, R, V)$ representing Alice’s initial belief state. An edge from a world v to a world w means $w \in R_a(v)$. We have $\mathcal{M} \models B_a \neg p \wedge \neg B_a p$. **Middle:** The standard event model $\mathcal{E} = (E, Q, pre)$ representing Bob’s announcement. The event e contains its own precondition $pre(e) = p$. The loop on e means $e \in Q_a(e)$. **Right:** The Kripke model $\mathcal{M} \otimes \mathcal{E} = (W^\mathcal{E}, R^\mathcal{E}, V^\mathcal{E})$ representing Alice’s belief state after the update, where $R_a^\mathcal{E}((v, e)) = \emptyset$ and so e.g. $\mathcal{M} \otimes \mathcal{E} \models B_a p \wedge B_a \neg p$.

3 Plausibility models

The “standard diagnosis” of the problem of modeling belief contraction in standard DEL is that we need a richer view of beliefs [12]. In particular, we need a framework where an agent believes not just any proposition entailed by her information, but those that are true at the ‘best’ or ‘most plausible’ possibilities compatible with her information [8, 12]. I believe that I am not dreaming. It is not *impossible* that I am, but those possibilities are less *plausible* than those in which I am awake. Similarly, one interpretation of Alice’s initial doxastic situation is that she believes that her office is closed, not because she has completely ruled out the possibility that it is open, but only that she judges that possibility to be relatively implausible. And it is this relative plausibility judgment that gets revised by Bob’s testimony.

This idea is formalized by modeling belief in terms of plausibility orderings over possible worlds. Formally, a *plausibility order on a set S* is a connected, reflexive and transitive relation $\leq \subseteq S \times S$ such that every non-empty subset $S' \subseteq S$ has at least one $\leq$ -minimal element, i.e. $Min_{\leq}(S') = \{w \in S' : \forall v \in S', w \leq v\} \neq \emptyset$.³ A *plausibility model* is a Kripke model $(W, \leq, V)$ with $\leq_a$ a plausibility order on W , for all agents $a \in Ag$. We write $\prec_a$ for its asymmetric component, i.e. $w \prec_a v$ if $w \leq_a v$ and $v \not\leq_a w$. The semantics of $\mathcal{L}$ over plausibility models is standard for propositional formulas, but belief is interpreted differently from above, capturing that the agent believes what is the case at the most plausible worlds:

$$\mathcal{M}, w \models B_a \varphi \text{ iff } \mathcal{M}, v \models \varphi \text{ for all } v \in Min_{\leq_a}(W).$$

A *plausibility event model* is an event model $\mathcal{E} = (E, \leq, pre)$ with $\leq_a$ a plausibility order on E for all $a \in Ag$. We write $\prec_a$ for its asymmetric component. As in standard DEL, event models update plausibility models via product update. The *plausibility-update* of $\mathcal{M} = (W, \leq, V)$ with $\mathcal{E}$ is the plausibility model $\mathcal{M} * \mathcal{E} = (W^\mathcal{E}, \leq^\mathcal{E}, V^\mathcal{E})$ where $W^\mathcal{E}$ and $V^\mathcal{E}$ are defined as in standard product updates and $\leq^\mathcal{E}$ is given

³Note that we assume the plausibility order to be *connected*. This is more restrictive than usual [8], but it is justified here as our framework only models belief and omits knowledge.

by $(w, e) \preceq_a^\mathcal{E} (v, f)$ iff $e <_a f$, or $e \equiv_a f$ and $w \preceq_a v$. Figure 2 shows how to model Example 2.1 via a plausibility update. Alice's change of mind is captured as a *soft* update where no world is ruled out as impossible, but their relative plausibility is switched.




Figure 2: **Left:** The plausibility model $\mathcal{M} = (W, \preceq, V)$ representing Alice's initial belief state. An edge from v to w means $w \preceq_a v$. We have $\mathcal{M} \models B_a \neg p \wedge \neg B_a p$. **Middle:** The plausibility event model $\mathcal{E} = (E, \leq, pre)$ capturing a soft update that p holds. An edge from f to e means $e \leq_a f$. **Right:** The plausibility model $\mathcal{M} * \mathcal{E}$ representing Alice's belief state after updating, with $\mathcal{M} * \mathcal{E} \models B_a p \wedge \neg B_a \neg p$.

3.1 Limitations of plausibility models

One limitation of the plausibility model of beliefs is that it validates axiom 4: $B_a \phi \rightarrow B_a B_a \phi$. For this reason, it cannot model agents who are mistaken about their own beliefs. Such cases, however, seem possible: an employer might profess that they believe in gender equality while consistently favoring one gender group in hiring. One way of describing this employer is that they in fact believe that candidates from one gender group are more qualified, even though they (falsely) believe that they don't believe that.

Another limitation of the approach, mentioned above, concerns belief contraction that does not involve adopting any new factual beliefs. A variant of Example 2.1 illustrates what we have in mind:

Example 3.1. As before, Alice believes that she locked the door to her office before she left. She runs into Bob, who tells her that *her office door might be open*. Given Bob's testimony, Alice comes to suspend judgment on whether she locked the door to her office before she left.

Let p denote the proposition *Alice's office door is open*. Alice's belief states before and after the update can be modeled by $\mathcal{M}$ and $\mathcal{N}$ in Figure 3, respectively.



Figure 3: **Left:** The plausibility model $\mathcal{M} = (W, \preceq, V)$ representing Alice's initial belief state, where $\mathcal{M} \models B_a \neg p \wedge \neg B_a p$. Conventions for edges are as in Fig. 2. **Right:** The plausibility model $\mathcal{N} = (W, \preceq^\mathcal{N}, V)$ representing Alice's belief state after updating on Bob's testimony, where $\mathcal{N} \models \neg B_a p \wedge \neg B_a \neg p$.

However, it turns out that this update *cannot* be modeled as a plausibility update. More precisely, let $\mathcal{K} \equiv \mathcal{K}'$ denote modal equivalence between two models $\mathcal{K}$ and $\mathcal{K}'$ with respect to $\mathcal{L}$.⁴ Then:

Theorem 3.2. *There is no plausibility event model $\mathcal{E}$ such that $\mathcal{M} * \mathcal{E} \equiv \mathcal{N}$.*

Proof. Let $\mathcal{M} = (W, \preceq, V)$ and $\mathcal{N} = (W, \preceq^\mathcal{N}, V)$ be as depicted in Figure 3. Suppose towards a contradiction that there is a plausibility event model $\mathcal{E} = \{E, \leq, pre\}$ such that $\mathcal{M} * \mathcal{E} = (W^\mathcal{E}, \preceq^\mathcal{E}, V^\mathcal{E})$ is a plausibility model that is modally equivalent to $\mathcal{N}$. We claim that, for any $(v, f) \in W^\mathcal{E}$, there exists a $(v, f') \neq (v, f)$ such that $(v, f') \prec_a^\mathcal{E} (v, f)$. It follows then that $\mathcal{M} * \mathcal{E}$ contains an infinite descending

⁴Recall that two Kripke (or plausibility) models $\mathcal{K} = (W, R, V)$ and $\mathcal{K}' = (W', R', V')$ are *modally equivalent* with respect to $\mathcal{L}$ if for every formula $\phi \in \mathcal{L}$ and every state $w \in W$, there exists a state $w' \in W'$ such that $\mathcal{K}, w \models \phi$ iff $\mathcal{K}', w' \models \phi$, and conversely, for every $w' \in W'$ there exists a state $w \in W$ with the same property.

chain, which contradicts the assumption that it is a plausibility model. Let $(v, f) \in W^{\mathcal{E}}$. By assumption, $(\mathcal{M} * \mathcal{E}, (v, f))$ is modally equivalent to $(\mathcal{N}, v)$. Since $(\mathcal{N}, v) \models \neg B_a p$, we have $(\mathcal{M} * \mathcal{E}, (v, f)) \models \neg B_a p$. So there exists $(w, e) \in W^{\mathcal{E}}$ such that $(w, e) \preceq_a^{\mathcal{E}} (v, f)$. For this (w, e) , it must be that $(\mathcal{M} * \mathcal{E}, (w, e))$ is modally equivalent to $(\mathcal{N}, w)$ (since it can't be modally equivalent to $(\mathcal{N}, v)$). Since $\mathcal{N}, w \models B_a p$, there exists $(v, f') \in W^{\mathcal{E}}$ such that $(v, f') \preceq_a^{\mathcal{E}} (w, e)$. Moreover, since $v \not\leq_a w$, we must have $f' <_a e \leq_a f$. So $f' <_a f$ and therefore $(v, f') \prec_a^{\mathcal{E}} (v, f)$. $\square$

Note that if two models are not modally equivalent with respect to $\mathcal{L}$, then they are not modally equivalent with respect to any extension of $\mathcal{L}$. So it follows from Theorem 3.2 that there is no plausibility event model $\mathcal{E}$ such that $\mathcal{M} * \mathcal{E}$ is modally equivalent to $\mathcal{N}$ with respect to $\mathcal{L}$ enriched with, e.g. modalities for conditional beliefs, safe beliefs or dynamic modalities. Relatedly, since modal equivalence is generally assumed to be a necessary condition for an adequate notion of bisimilarity, it also follows that there is no plausibility event model $\mathcal{E}$ such that $\mathcal{M} * \mathcal{E}$ is bisimilar to $\mathcal{N}$ with respect to any adequate notion of bisimilarity (see e.g. [16, 2, 3] for different notions of bisimilarity for plausibility structures).

4 A new model for belief contraction

Theorem 3.2 shows that plausibility updates can't weaken an agent's belief in a way that makes her consider the two worlds equiplausible. We now introduce a framework for belief contraction that can achieve that. In general, belief contraction can be caused by many different kinds of epistemic events. In this section, we focus on a simple case of belief contraction due to *hedged public announcements*, like Bob's testimony that Alice's door *might* be open.⁵ Formally, we represent the effects of such announcements by the dynamic modality $[\div \varphi] \psi$. Intuitively, it says that ψ is the case after the hedged public announcement that φ *might be false*. Let the *contraction language* $\mathcal{L}^{\div}$ be the language generated by the grammar of $\mathcal{L}$ extended with the following clauses for the dynamic modality and a universal modality: $\varphi := [\div \varphi] \varphi \mid \forall \varphi$. We define $\exists \varphi := \neg \forall \neg \varphi$.

A standard (non-hedged) public announcement that φ *is false* is modeled as eliminating all the φ -possibilities from all agents' doxastic spaces [29]. A natural thought, then, is that a hedged public announcement that φ *might be false* should add all the $\neg \varphi$ -possibilities to all agents' doxastic spaces. However, this is not quite right. Adding all such possibilities to all agents' doxastic spaces may force those who already consider $\neg \varphi$ possible to update their beliefs and start considering additional $\neg \varphi$ -possibilities, thereby losing more beliefs than is warranted by the announcement. This suggests that, unlike standard public announcements, the epistemic effects of a hedged announcement that φ *might be false* depend on the agents' initial beliefs, and in particular on whether they believe φ prior to the announcement. For this section, we assume that this is the only relevant difference: if an agent already considers $\neg \varphi$ possible, then the announcement φ *might be false* has no effects on her beliefs; if the agent believes φ , then the announcement leads her to consider all $\neg \varphi$ -worlds as possible.

Definition 4.1 (Contraction on φ). Given a Kripke model $\mathcal{M} = (W, R, V)$ and a formula $\varphi \in \mathcal{L}^{\div}$. Let $\mathcal{M}^{\div \varphi} = (W, R^{\div \varphi}, V)$ be defined by, for all $a \in Ag$ and all $w \in W$:

$$R_a^{\div \varphi}(w) = \begin{cases} R_a(w) \cup \{v \in W : \mathcal{M}, v \models \neg \varphi\} & \text{if } \mathcal{M}, w \models B_a \varphi \\ R_a(w) & \text{otherwise} \end{cases}$$

⁵The name *hedged public announcement* reflects the public and non-committal ("hedged") character of the announcement, which opens up possibilities rather than restricting them—much like epistemic modals in so-called "hedged assertions" [13].

Call $\mathcal{M}^{\div \varphi}$ the model obtained from $\mathcal{M}$ by contracting on φ . Note that it is always the case that $R_a(w) \subseteq R_a^{\div \varphi}(w)$. This means that a hedged public announcement always *extends* the initial Kripke model. The semantics of $\mathcal{L}^{\div}$ is given by the semantics of $\mathcal{L}$ over Kripke models extended with these clauses:

$$\begin{aligned} \mathcal{M}, w \models [\div \varphi] \psi & \text{ iff } \mathcal{M}^{\div \varphi}, w \models \psi; \\ \mathcal{M}, w \models \forall \varphi & \text{ iff for all } v \in W, \mathcal{M}, v \models \varphi. \end{aligned}$$

Consider again Example 3.1. In the present framework, the effect of Bob's announcement that *Alice's office door might be open* can be modeled by contracting on the proposition *Alice's office door is closed* ($\neg p$), which amounts to adding world v to Alice's set of doxastic possibilities at w and v . See Figure 4.



Figure 4: **Left:** The Kripke model $\mathcal{M}$ (same as Figure 1), representing Alice's initial belief state. **Right:** The Kripke model $\mathcal{M}^{\div \neg p}$ representing Alice's belief state after updating on Bob's hedged announcement, obtained by contracting on $\neg p$. We have $\mathcal{M}^{\div \neg p} \models \neg B_a \neg p \wedge \neg B_a p$, so Alice lost a belief in $\neg p$.

One natural question is what is the logic of the contraction update as defined above. The next theorem shows that such logic is given by the system presented in Table 1.

CL	all classical propositional tautologies
$\forall K$	$\forall(\varphi \rightarrow \psi) \rightarrow (\forall \varphi \rightarrow \forall \psi)$
$\forall T$	$\forall \varphi \rightarrow \varphi$
$\forall 5$	$\neg \forall \varphi \rightarrow \forall \neg \forall \varphi$
$\forall B$	$\forall \varphi \rightarrow B_a \varphi$
BK	$B_a(\varphi \rightarrow \psi) \rightarrow (B_a \varphi \rightarrow B_a \psi)$
A1	$[\div \varphi] p \leftrightarrow p$
A2	$[\div \varphi] \neg \psi \leftrightarrow \neg [\div \varphi] \psi$
A3	$[\div \varphi] (\psi \wedge \chi) \leftrightarrow [\div \varphi] \psi \wedge [\div \varphi] \chi$
A4	$[\div \varphi] \forall \psi \leftrightarrow \forall [\div \varphi] \psi$
A5	$[\div \varphi] B_a \psi \leftrightarrow (B_a [\div \varphi] \psi \wedge (B_a \varphi \rightarrow \forall (\neg \varphi \rightarrow [\div \varphi] \psi)))$
MP	From φ and $\varphi \rightarrow \psi$, infer ψ
RE	From $\varphi \leftrightarrow \psi$, infer $\chi[\varphi/p] \leftrightarrow \chi[\psi/p]$
NEC	Necessitation rules for all modalities ⁶

Table 1: The Theory of Hedged Public Announcement Logic (HPAL)

Theorem 4.2. *The dynamic logic of contraction updates is completely axiomatized by HPAL.*

Proof sketch. Soundness follows by straightforward validity arguments. For completeness, given the reduction axioms, every formula with dynamic modalities is provably equivalent in HPAL to a formula without. So the completeness of HPAL follows from the completeness of the basic doxastic logic with the universal modality [14, Theorem 7.2]. $\square$

Besides axioms and rules of inference of the basic modal logic K with the universal modality (which is an S5 modality), the theory contains inference rules and reduction axioms for the contraction modality that

⁶Notice that necessitation for the belief modality is redundant, as it is derivable from $\forall B$ and $\forall$ -necessitation.

parallel the reduction axioms in Public Announcement Logic (PAL) [29, 20, 31]. Axiom A5 is specific to our logic of contraction. It says that after contraction on φ (viz. a hedged announcement that φ might be false), an agent believes ψ iff she doesn't believe φ prior to the announcement and she believes that ψ holds after the announcement, or she believes φ prior to the announcement and at any $\neg\varphi$ -worlds, ψ holds after the announcement.

5 Preservation and Moorean formulas

A well-known phenomenon in public announcement logic is that some formulas become false after they are announced. As a result, the formula is not believed after the announcement. In PAL, formulas that become false or are not believed after an announcement are both called *unsuccessful formulas*. Let $[\!|\varphi|\!]$ be the dynamic modality “after the announcement of φ ” in PAL.⁷ There are thus two closely related notions of *success* in PAL: (i) what is announced is true after its announcement ($[\!|\varphi|\!]\varphi$ is valid); (ii) what is announced is believed after its announcement ($[\!|\varphi|\!]B_a\varphi$ is valid for all $a \in Ag$). For clarity, we will say φ is *preserved in PAL* if it satisfies (i); φ is *successful in PAL* if it satisfies (ii). It is well-known that, in PAL, any formula that satisfies (i) also satisfies (ii), though the converse does not hold.⁸ In this section, we focus on developing a notion of preservation for HPAL that is analogous to (i) for PAL.

A natural starting point is to say that φ is *preserved* if it is true after the hedged announcement that φ might be true, viz. $[\div\neg\varphi]\varphi$ is valid. But this isn't quite right: while a public announcement that φ is true *eliminates* all $\neg\varphi$ -possibilities, a hedged public announcement that φ *might* be true doesn't eliminate any possibilities, and so any possibility where φ is false before the announcement remains in the model (though they may satisfy φ in the new model after the announcement). This observation suggests the following alternative definition of preservation for hedged public announcements (let $\mathcal{L}$ be a class of models that validates logic **L**):

Definition 5.1 (Preservation). φ is *preserved* (in the logic **L**) if $\models_{\mathcal{L}} \varphi \rightarrow [\div\neg\varphi]\varphi$

Not all formulas in $\mathcal{L}^{\div}$ are preserved in HPAL. Consider Example 3.1 again. Suppose Bob tells Alice instead that *she might have a false belief that her office door is closed*, viz. it might be true that $p \wedge B_a\neg p$. Figure 5 shows that this hedged announcement is not preserved.



Figure 5: Alice's belief update as a contraction on $\neg(p \wedge B_a\neg p)$. The formula $p \wedge B_a\neg p$ is true at v before the update (on the left) but false at v after the update (on the right). Hence, $p \wedge B_a\neg p$ is not preserved.

More generally, say φ is *self-refuting* (in the logic **L**) if $\models_{\mathcal{L}} \varphi \rightarrow [\div\neg\varphi]\neg\varphi$. Clearly, φ is not preserved if it is self-refuting. The above example is a special case of the following result:

⁷In PAL, the formula $[\!|\varphi|\!]\psi$ intuitively says ψ is true after eliminating all $\neg\varphi$ -possibilities from the model; see [20, Chapter 4] and references therein for more on PAL.

⁸The proof that (i) implies (ii) is analogous to the “only if”-direction of [20, Proposition 4.32], though here the modality is belief rather than knowledge, and in particular it does not necessarily satisfy the T-axiom, which is the reason why the other direction does not hold in general. For example, suppose $Ag = \{a\}$, let $\varphi := p \wedge B_a\neg p \wedge \neg B_a\perp$, and consider a Kripke model and a world w in it where such φ is true. Then, after the public announcement of φ , which eliminates all possibilities where φ is not true and so also all worlds that were accessible for a at w , a will have inconsistent beliefs at w and believe (trivially) that φ is true. Hence, $[\!|\varphi|\!]B_a\varphi$ is valid. On the other hand, $[\!|\varphi|\!]\varphi$ is invalid and in particular $[\!|\varphi|\!]\neg B_a\perp$ is invalid.

Proposition 5.2. *The formulas $p \wedge \exists B_a \neg p$ and $p \wedge B_a \neg p$ are self-refuting in HPAL.*⁹

Proof. Let $\varphi := p \wedge \exists B_a \neg p$. Suppose $\mathcal{M}, w \models p \wedge \exists B_a \neg p$. Then $\mathcal{M} \models \exists B_a \neg p$ and so $\mathcal{M}^{\div \neg \varphi} = \mathcal{M}^{\div \neg p}$. Thus $\mathcal{M}^{\div \neg \varphi} \models \forall \tilde{B}_a p$ and so $\mathcal{M}^{\div \neg \varphi}, w \models \neg \varphi$. Similarly, let $\psi := p \wedge B_a \neg p$. Suppose $\mathcal{M}, w \models \psi$. Since $B_a \neg p \models B_a \neg \psi$, $\mathcal{M}, w \models B_a \neg \psi$. So $w \in R_a^{\div \neg \psi}(w)$. Since $\mathcal{M}, w \models p$, $\mathcal{M}^{\div \neg \psi}, w \models \tilde{B}_a p$. Given that $\tilde{B}_a p \models \neg \psi$, $\mathcal{M}^{\div \neg \psi}, w \models \neg \psi$. $\square$

So which formulas are preserved? Recall that a formula φ is *existential* if it is built only from literals, $\wedge$, $\vee$ and $\tilde{B}_a$; φ is *universal* if it is built only from literals, $\wedge$, $\vee$ and B_a . It's well-known that universal formulas remain true after being (standardly) publicly announced, in any normal modal logic [20]. Their dual—existential formulas—are always preserved by hedged public announcements:

Theorem 5.3. *If $\varphi \in \mathcal{L}^{\div}$ is equivalent in HPAL to an existential formula, then it is preserved in HPAL.*

Proof. Suppose φ is equivalent in HPAL to an existential formula. Then φ is preserved under model extension [4]. Since $\mathcal{M}^{\div \neg \varphi}$ extends $\mathcal{M}$, if $\mathcal{M}, w \models \varphi$, then $\mathcal{M}^{\div \neg \varphi}, w \models \varphi$, i.e. $\mathcal{M}, w \models \varphi \rightarrow [\div \neg \varphi] \varphi$. $\square$

One upshot of Theorem 5.3 is that, while the “strong” Moorean sentence $p \wedge B_a \neg p$ is never preserved, the “weak” Moorean sentence $p \wedge \neg B_a p$, which is equivalent to an existential formula, is preserved.

So being logically equivalent to an existential formula is sufficient for being preserved. But it is not necessary. In particular, the non-existential formula $p \wedge B_a p$ (agent a has a true belief that p) is preserved.

Proposition 5.4. *The formula $p \wedge B_a p$ is preserved in HPAL.*

Proof. Let $\mathcal{M} = (W, R, V)$ be a Kripke model with $\mathcal{M}, w \models p \wedge B_a p$. As $R_a^{\div \neg (p \wedge B_a p)}(w) \subseteq R_a(w) \cup \{v \in W : \mathcal{M}, v \models p \wedge B_a p\} \subseteq \{v \in W : \mathcal{M}, v \models p\}$ and p is preserved, then $\mathcal{M}^{\div \neg (p \wedge B_a p)}, w \models p \wedge B_a p$. $\square$

It would be nice to have a complete characterization of the set of preserved formulas, which we leave as an open question. But there is a partial result. Let's consider *knowledge* instead of *belief* for a moment, and restrict our attention to single-agent models. The standard logic for knowledge is **S5**, and extending HPAL with **S5** axioms implies that all formulas in $\mathcal{L}$ are preserved.

Theorem 5.5. *Suppose $|Ag| = 1$. Then every formula in $\mathcal{L}$ is preserved in HPAL extended with **S5**.*

Proof sketch. In **S5**, if $\mathcal{M}, w \models \varphi$ with $\varphi \in \mathcal{L}$, then every world in the equivalence class of w satisfies $\tilde{B}_a \varphi$ and so contracting on $\neg \varphi$ does not change relations inside those classes. Then, $(\mathcal{M}, w)$ and $(\mathcal{M}^{\div \neg \varphi}, w)$ are bisimilar, and since $\mathcal{L}$ is bisimulation-invariant, φ is preserved. $\square$

Note that model $\mathcal{M}$ in Figure 5 validates **KD45**, so the preservation result doesn't hold in these settings.¹⁰

6 AGM-style contraction principles

In the AGM literature, belief contraction is modeled as a static operation on the set of formulas, representing the agent's belief base, and is characterized by a set of axioms. As [9] points out, many of those axioms do not naturally hold in the dynamic setting where belief contraction induces model transformations. In this section, we give an overview of which of the AGM axioms hold/fail for our contraction operator. Following [30], we render AGM statements that express that the agent believes χ after contracting by φ , by the formula $[\div \varphi] B_a \chi$. The modal version of the AGM axioms for contraction can then be expressed by the following formulas, where $\varphi, \psi, \chi \in \mathcal{L}^{\div}$:

⁹Note that the Moorean sentence $p \wedge B_a \neg p$ is unsatisfiable in **S5**, while $p \wedge \exists B_a \neg p$ is satisfiable even if belief is factive.

¹⁰As an anonymous referee pointed out, the result does not generalize to the multi-agent setting.

- **Closure.** $[\div \varphi](B_a(\psi \rightarrow \chi) \rightarrow (B_a\psi \rightarrow B_a\chi))$.
- **Success.** $\exists \neg \varphi \rightarrow [\div \varphi] \neg B_a\varphi$.
- **Inclusion.** $[\div \varphi]B_a\psi \rightarrow B_a\psi$.
- **Vacuity.** $\neg B_a\varphi \rightarrow (\chi \leftrightarrow [\div \varphi]\chi)$.
- **Extensionality.** If $\varphi \leftrightarrow \psi$ then $\models [\div \varphi]\chi \leftrightarrow [\div \psi]\chi$.
- **Conjunctive Inclusion.** $[\div(\varphi \wedge \psi)] \neg B_a\varphi \rightarrow ([\div(\varphi \wedge \psi)]B_a\chi \rightarrow [\div \varphi]B_a\chi)$.
- **Conjunctive Overlap.** $([\div \varphi]B_a\chi \wedge [\div \psi]B_a\chi) \rightarrow [\div(\varphi \wedge \psi)]B_a\chi$.
- **Consistency.** $\neg B_a\perp \rightarrow [\div \varphi] \neg B_a\perp$.

From the list above we omitted the axiom called **Recovery** (expansion by φ after contraction by φ restores the agent's original beliefs), as it involves an expansion operation which we do not focus on.¹¹ We added to this list a natural property called **Consistency**, which is discussed in [30, 26].

Some of these properties are valid in HPAL. In particular, **Closure** and **Extensionality** hold unrestrictedly in our semantic framework. **Consistency** also holds, as the contraction operation never removes edges. For the other properties, the situation is more mixed, as the next two subsections show.

6.1 Success

Recall that φ is successful in PAL if it is believed after it is announced: $[\div \varphi]B_a\varphi$ is valid for all $a \in Ag$. A natural analogue of this notion in our contraction settings is that φ is *not believed* after a hedged announcement that it might be false, viz. $[\div \varphi] \neg B_a\varphi$ for all $a \in Ag$. The only caveat concerns the case where φ is valid in the model. In that case, since there are no $\neg\varphi$ -possibilities, contracting on φ leaves the model unchanged and so may leave agents' beliefs in φ unchanged.¹² This illustrates why **Success** is given by the conditional $\exists \neg \varphi \rightarrow [\div \varphi] \neg B_a\varphi$. Say that a formula φ is *successful* if it satisfies this property.

As with preservation, not all formulas are successful. For instance, the hedged announcement that $p \wedge B_a \neg p$ might be true in Figure 5 is unsuccessful: after updating on Bob's announcement that she might have a false belief that $\neg p$, by contracting on $\neg(p \wedge B_a \neg p)$, Alice will come to believe that she believes neither p nor $\neg p$, and so she will believe that she doesn't have a false belief that p (i.e. $\mathcal{M}^{\div \neg(p \wedge B_a \neg p)} \models B_a \neg(p \wedge B_a \neg p)$). In PAL, if a formula remains true after being announced, then it is believed after the announcement [20, Proposition 4.32]. A similar result holds for HPAL.

Proposition 6.1. *In HPAL, if $\neg\varphi \in \mathcal{L}^{\div}$ is preserved, then φ is successful.*

Proof. Let $\mathcal{M} = (W, R, V)$ and suppose $\mathcal{M}, w \models \exists \neg \varphi$. Assume that $\neg\varphi$ is preserved. We show that $\mathcal{M}, w \models [\div \varphi] \neg B_a\varphi$. There are two cases. First, suppose that $\mathcal{M}, w \models \neg B_a\varphi$. Then, there is $v \in R_a(w) = R_a^{\div \varphi}(w)$ such that $\mathcal{M}, v \models \neg\varphi$. As $\neg\varphi$ is preserved, $\mathcal{M}, v \models [\div \varphi] \neg\varphi$ and so $\mathcal{M}, w \models [\div \varphi] \neg B_a\varphi$. Second, suppose $\mathcal{M}, w \models B_a\varphi$. Since $\mathcal{M}, w \models \exists \neg \varphi$, there is a $v \in W$ with $\mathcal{M}, v \models \neg\varphi$ and $v \in R_a^{\div \varphi}(w)$. By preservation of $\neg\varphi$, $\mathcal{M}, v \models [\div \varphi] \neg\varphi$ and so $\mathcal{M}, w \models [\div \varphi] \neg B_a\varphi$. $\square$

Given that existential formulas are preserved (Theorem 5.3), we then obtain the following corollary, which gives a sufficient condition for a formula to be successful:

¹¹Notice, however, that **Recovery** does not hold in general in this framework: while the contraction operator $[\div \varphi]$ can informally be viewed as a dual of the public announcement operator $[\div \varphi]$, it is not a formal dual, in the sense that contraction by φ followed by expansion by φ does not in general recover the original model (whether we model expansion by world- or arrow-eliminations). We explore this issue in a longer version of the paper.

¹²This mirrors the AGM constraint that beliefs in validities cannot be given up [1, 25].

Corollary 6.2. *If $\varphi \in \mathcal{L}^\div$ is equivalent in HPAL to a universal formula, then φ is successful in HPAL.*

Proof. If φ is equivalent to a universal formula, $\neg\varphi$ is equivalent to an existential formula. By Theorem 5.3 existential formulas are preserved, and by Prop. 6.1 if $\neg\varphi$ is preserved, then φ is successful. $\square$

6.2 Other properties

The previous section showed that **Success** is not valid in HPAL. However, by Corollary 6.2, all universal formulas are successful and in particular all propositional formulas are successful. The same holds for **Inclusion**, **Conjunctive Inclusion** and **Conjunctive Overlap**. These postulates fail for reasons similar to the failure of **Success**. For instance, in Figure 4, after Bob’s hedged announcement, Alice comes to believe that she doesn’t believe p —a belief that she didn’t have before, which is a failure of **Inclusion** ($\mathcal{M}, w \models \neg B_a \neg B_a \neg p \wedge [\div \neg p] B_a \neg B_a \neg p$). This is to be expected: on the dynamic approach, the truths of doxastic formulas can be influenced by hedged announcements. However, all three postulates hold if the believed formula is propositional. Let $\mathfrak{K}$ be the class of all Kripke models.

Proposition 6.3. *The following are valid in HPAL, where χ and λ are propositional and $\varphi, \psi \in \mathcal{L}^\div$:*

1. **Propositional Inclusion.** $[\div \varphi] B_a \chi \rightarrow B_a \chi$.
2. **Propositional Conjunctive Inclusion.** $[\div \lambda \wedge \psi] \neg B_a \lambda \rightarrow ([\div \lambda \wedge \psi] B_a \chi \rightarrow [\div \lambda] B_a \chi)$.
3. **Propositional Conjunctive Overlap.** $([\div \varphi] B_a \chi \wedge [\div \psi] B_a \chi) \rightarrow [\div \varphi \wedge \psi] B_a \chi$.

Proof sketch. **Propositional Inclusion** follows from $R_a(w) \subseteq R_a^{\div \varphi}(w)$ and from propositional truth being unaffected by contraction. **Propositional Conjunctive Inclusion** and **Overlap** follow because contracting on a conjunction introduces no alternatives beyond those required to give up one of its conjuncts. $\square$

Finally, while **Vacuity** is also among the properties that do not hold unrestrictedly, the following strong version holds:

- **Strong Vacuity.** $\bigwedge_{a \in Ag} \forall \neg B_a \varphi \rightarrow (\chi \leftrightarrow [\div \varphi] \chi)$.

This principle strengthens the standard **Vacuity** principle listed above by requiring that *no agent* believes the contracted formula *at any world in the model*. Such restrictions reflect the multi-agent and public nature of our contraction operation. They are not required in classic AGM theory, where beliefs are propositional, nor in dynamic doxastic logic, which is single-agent [1, 30].

Proposition 6.4. *Strong Vacuity is a theorem of HPAL.*

Proof. Let $\varphi, \chi \in \mathcal{L}^\div$ and consider a Kripke model $\mathcal{M} = (W, R, V)$ and a world w in it. Assume that $\mathcal{M}, w \models \forall \neg B_a \varphi$, for all $a \in Ag$, i.e., $\mathcal{M}, u \models \neg B_a \varphi$ for all $u \in W$ and $a \in Ag$. Then $R_a(u) = R_a^{\div \varphi}(u)$ for all worlds u and all agents a , i.e., $\mathcal{M} = \mathcal{M}^{\div \varphi}$ and so $\mathcal{M}, w \models \chi$ iff $\mathcal{M}^{\div \varphi}, w \models \chi$, that is, $\mathcal{M}, w \models \chi$ iff $\mathcal{M}, w \models [\div \varphi] \chi$. $\square$

7 Generalized DEL

In this section, we introduce a generalized DEL framework that can model belief contraction resulting from different kinds of announcements, like hedged private announcements.

Let the *DEL language* $\mathcal{L}_{DEL}$ be the language generated by the grammar of the doxastic language $\mathcal{L}$ extended with the following clauses $\varphi ::= [\mathcal{E}, e] \varphi \mid \forall \varphi$, where $\mathcal{E}$ is a generalized event model (introduced below) and e is an event in it. The formula $[\mathcal{E}, e] \varphi$ reads “after $(\mathcal{E}, e)$ happens, φ is the case”.

Definition 7.1 (Generalized event model). A *generalized event model* is a tuple $\mathcal{E} = (E, Q, Q^+, pre)$ where E and Q_a are defined as in standard event models, $Q_a^+ \subseteq E \times E$ is an accessibility relation between events defined for all agents $a \in Ag$, and $pre : E \rightarrow \mathcal{L}_{DEL}$ is a precondition function.

A generalized event model is a standard event model extended with additional accessibility relations Q_a^+ , for each agent $a \in Ag$, and in which preconditions can be formulas of the dynamic language $\mathcal{L}_{DEL}$. As before, we use notation $Q_a^+(e) = \{f \in E : (e, f) \in Q_a^+\}$. Intuitively, both $Q_a(e)$ and $Q_a^+(e)$ represent events that are accessible, and thus considered possible, by agent a at event e . The difference, which will become formally clear with the product update definition below, is that $Q_a^+(e)$ contains alternatives that are *introduced by event e as possible for agent a* (irrespective of a 's prior beliefs), while $Q_a(e)$ is the standard accessibility relation containing possibilities that are not *ruled out for agent a by event e* .

The distinction is illustrated by the event model $\mathcal{E}$ in Figure 6, which is the event model corresponding to Example 3.1. Intuitively, f represents the event *Bob tells Alice her office door might be open and it is actually closed*, and h represents the event *Bob tells Alice her office door might be open and it is actually open*. For Alice, Bob's announcement doesn't rule out either f or h , but it introduces h as possible. Formally, this means $Q_a(f) = Q_a(h) = \{f, h\}$, while $Q_a^+(f) = Q_a^+(h) = \{h\}$.




Figure 6: **Left:** The Kripke model $\mathcal{M}$ representing Alice's initial belief states. **Middle:** The generalized event model $\mathcal{E} = (E, Q, Q^+, pre)$ representing Bob's announcement to Alice that p might hold. An edge from an event f to an event h labeled by $+a$ means that $h \in Q_a^+(f)$. We omit the Q -edges since $Q_a(e) = E$ for all $e \in E$. **Right:** The Kripke model $\mathcal{M} \otimes \mathcal{E} = (W^\mathcal{E}, R^\mathcal{E}, V^\mathcal{E})$ representing Alice's belief states after Bob's announcement to Alice.

Given this interpretation, the product update is then defined as the following:

Definition 7.2 (Generalized product update). Given a Kripke model $\mathcal{M} = (W, R, V)$ and a generalized event model $\mathcal{E} = (E, Q, Q^+, pre)$, their product update $\mathcal{M} \otimes \mathcal{E}$ is defined as $\mathcal{M} \otimes \mathcal{E} = (W^\mathcal{E}, R^\mathcal{E}, V^\mathcal{E})$ where $W^\mathcal{E}$ and $V^\mathcal{E}$ are defined as in standard product updates (cf. Section 2) and $R^\mathcal{E}$ is given by:

$$(v, f) \in R_a^\mathcal{E}(w, e) \text{ iff either } f \in Q_a^+(e), \text{ or } v \in R_a(w) \text{ and } f \in Q_a(e).$$

This product update behaves like a standard product update in eliminating possibilities (cf. Section 2), but additionally allows to expand the set of accessible worlds by adding all worlds (v, f) where f is a newly considered possibility. In particular, the first disjunct allows agents to start considering worlds paired with events in Q_a^+ , regardless of the original accessibility relation.

The semantics of the dynamic modality is standard [20]:

$$\mathcal{M}, w \models [\mathcal{E}, e]\varphi \quad \text{iff} \quad \text{if } \mathcal{M}, w \models pre(e) \text{ then } \mathcal{M} \otimes \mathcal{E}, (w, e) \models \varphi.$$

The following definition shows that generalized DEL can express contraction as induced by hedged public announcements.

Definition 7.3 (Event model for contraction on φ). Let $\varphi \in \mathcal{L}_{DEL}$. An *event model for contraction on φ* is a generalized event model $\mathcal{E}(\div \varphi) = (E, Q, Q^+, pre)$ where

1. $E = \{\varphi \wedge \bigwedge_{a \in A} \neg B_a \varphi \wedge \bigwedge_{a \in Ag \setminus A} B_a \varphi : A \subseteq Ag\} \cup \{\neg \varphi \wedge \bigwedge_{a \in A} \neg B_a \varphi \wedge \bigwedge_{a \in Ag \setminus A} B_a \varphi : A \subseteq Ag\}.$

2. For all $a \in \text{Ag}$ and $e \in E$, $Q_a(e) = E$ and $Q_a^+(e) = \begin{cases} \{f \in E : \text{pre}(f) \models \neg\varphi\} & \text{if } \text{pre}(e) \models B_a\varphi; \\ \emptyset & \text{otherwise.} \end{cases}$
3. For all $e \in E$, $\text{pre}(e) = e$.

Event model $\mathcal{E}(\div\varphi)$ contains events for all possible configurations of the truth value of φ and which agents believe φ . The relation Q_a is universal, capturing that the announcement reveals nothing about which event occurred. The relation Q_a^+ captures contraction: agents who believed φ come to consider $\neg\varphi$ -events possible, while the others are unaffected.

Proposition 7.4. For any Kripke model $\mathcal{M}$ and formula $\varphi \in \mathcal{L}$, $\mathcal{M}^{\div\varphi}$ is isomorphic to $\mathcal{M} \otimes \mathcal{E}(\div\varphi)$.

Proof sketch. Both updates preserve all worlds of $\mathcal{M}$, so there is an isomorphism between the updated models: as $R_a(w) \subseteq R_a^{\div\varphi}(w)$ and $Q_a(e) = E$ for all $a \in \text{Ag}, e \in E$, all relations in $\mathcal{M}$ are preserved by both updates. Also, the only relations added in both cases are from $B_a\varphi$ -worlds to $\neg\varphi$ -worlds. $\square$

Next we consider an example involving hedged private announcement:

Example 7.5. Alice just told Clark that she locked her office door before she left. Later, while Clark is evidently distracted, Bob privately tells Alice that her office door might actually be open. Alice comes to suspend judgment on whether she locked the door, while Clark does not notice this exchange, continuing to believe that Alice's office door is closed and that Alice believes that.

We model Bob's private announcement using the generalized event model $\mathcal{E}$ as described below:




Figure 7: **Left:** The Kripke model $\mathcal{M}$ representing Alice and Clark's initial belief states. We have $\mathcal{M} \models B_a\neg p \wedge B_c\neg p \wedge B_cB_a\neg p$. **Middle:** The generalized event model $\mathcal{E} = (E, Q, Q^+, \text{pre})$ representing Bob's private announcement to Alice that p might hold. An edge from an event f to an event h labeled by $+a$ means that $h \in Q_a^+(f)$. If an edge's label does not contain $+$, then the edge is a Q -relation, e.g. $e \in Q_c(f)$. **Right:** The Kripke model $\mathcal{M} \otimes \mathcal{E} = (W^\otimes, R^\otimes, V^\otimes)$ representing Alice and Clark's belief states after Bob's private announcement to Alice.

We now look into what is the logic of generalized updates. First, we define the composition of two generalized event models and then show that it is equivalent to the sequential updates of the two.

Definition 7.6 (Composition). Given two generalized event models $\mathcal{E} = (E^\mathcal{E}, Q^\mathcal{E}, Q^{+\mathcal{E}}, \text{pre}^\mathcal{E})$ and $\mathcal{F} = (E^\mathcal{F}, Q^\mathcal{F}, Q^{+\mathcal{F}}, \text{pre}^\mathcal{F})$, then their composition is $\mathcal{E} \circ \mathcal{F} = (E, Q, Q^+, \text{pre})$ where

- $E = E^\mathcal{E} \times E^\mathcal{F}$;
- $(e', f') \in Q_a((e, f))$ iff $e' \in Q_a^\mathcal{E}(e)$ and $f' \in Q_a^\mathcal{F}(f)$;
- $(e', f') \in Q_a^+((e, f))$ iff $f' \in Q_a^{+\mathcal{F}}(f)$, or $f' \in Q_a^\mathcal{F}(f)$ and $e' \in Q_a^{+\mathcal{E}}(e)$;
- $\text{pre}((e, f)) = \text{pre}^\mathcal{E}(e) \wedge [\mathcal{E}, e] \text{pre}^\mathcal{F}(f)$.

CL	all classical propositional tautologies
$\forall K$	$\forall(\varphi \rightarrow \psi) \rightarrow (\forall\varphi \rightarrow \forall\psi)$
$\forall T$	$\forall\varphi \rightarrow \varphi$
$\forall 5$	$\neg\forall\varphi \rightarrow \forall\neg\varphi$
$\forall B$	$\forall\varphi \rightarrow B_a\varphi$
BK	$B_a(\varphi \rightarrow \psi) \rightarrow (B_a\varphi \rightarrow B_a\psi)$
A1	$[\mathcal{E}, e]p \leftrightarrow (pre(e) \rightarrow p)$
A2	$[\mathcal{E}, e]\neg\varphi \leftrightarrow (pre(e) \rightarrow \neg[\mathcal{E}, e]\varphi)$
A3	$[\mathcal{E}, e](\varphi \wedge \psi) \leftrightarrow ([\mathcal{E}, e]\varphi \wedge [\mathcal{E}, e]\psi)$
A4	$[\mathcal{E}, e]\forall\varphi \leftrightarrow (pre(e) \rightarrow \bigwedge_{f \in E} \forall[\mathcal{E}, f]\varphi)$
A5	$[\mathcal{E}, e]B_a\varphi \leftrightarrow (pre(e) \rightarrow (\bigwedge_{f \in Q_a^+(e)} \forall[\mathcal{E}, f]\varphi \wedge \bigwedge_{f \in Q_a(e)} B_a[\mathcal{E}, f]\varphi))$
A6	$[\mathcal{E}, e][\mathcal{F}, f]\varphi \leftrightarrow [\mathcal{E} \circ \mathcal{F}, (e, f)]\varphi$
MP and necessitation rules for all modalities	

Table 2: The Theory GDEL

The following shows that this is given by the system presented in Table 2.¹³

Theorem 7.7. *The dynamic logic of generalized updates is completely axiomatized by GDEL.*

Proof sketch. The proof strategy is analogous to that used for Theorem 4.2, using the composition axiom A6 instead of the inference rule RE. $\square$

The theory GDEL contains the static axioms of HPAL, together with the inference rules for all modalities. Additionally, it contains reduction axioms for event models analogous to those of standard DEL. Axiom A5 captures the effect of generalized updates on belief via the relations Q_a and Q_a^+ .

The next theorem shows that generalized updates always produce a model that is a simulation of a refinement of the initial model (refinement and simulation are defined standardly [14], but see Appendix). As noted in [15, 18], refinements decrease uncertainty by eliminating possibilities, while simulations capture increases in uncertainty by adding possibilities. Generalized updates thus can be seen as updates where agents may rule out possibilities while starting to consider new ones, as in belief revision.

Theorem 7.8. *For any Kripke model $\mathcal{M}$ and generalized event model $\mathcal{E}$, $\mathcal{M} \otimes \mathcal{E}$ is a simulation of a refinement of $\mathcal{M}$.*

Proof sketch. The idea is to split a generalized update into two steps. First ignore the Q_a^+ edges and obtain a standard DEL update, which gives a refinement of the initial model [15]. The full generalized update then only adds edges to this model, so it is a simulation of that refinement. $\square$

Before concluding this section, we discuss one limitation of the contraction operation as defined in Section 4, and how the generalized product update helps to overcome it.

Example 7.9. Alice and Clark walk down the hallway. Alice believes that she closed her office door ($\neg p$) and that there will be a department meeting in the afternoon (q). Clark, however, is uncertain about both questions, though he believes that Alice believes the truth no matter what it is, and so does Alice. They then encounter Bob, who tells them that Alice's office door might be open.

¹³Note that, unlike the axiomatization of HPAL, the axiomatization of GDEL includes a composition axiom (A6) but without RE as a valid rule of inference. We conjecture that we can omit A6 and include RE as a valid inference rule, though we leave the verification of this conjecture to future work.

Intuitively, given Bob's announcement, Clark will come to think that, if Alice initially believes that her office door is closed, then she'll come to suspend judgment about whether her office door is open, but she will remain confident about whether there will be a department meeting in the afternoon. So Alice's and Clark's beliefs should be modeled by $\mathcal{N}$ in Figure 8.




Figure 8: **Left:** The Kripke model $\mathcal{M}$ representing Alice and Clark's initial belief states. We have $\mathcal{M} \models B_c(B_ap \vee B_a\neg p) \wedge B_c(B_aq \vee B_a\neg q)$. **Right:** The Kripke model $\mathcal{N}$ representing Alice and Clark's belief states after Bob's hedged announcement, where $\mathcal{N}, w \models \neg B_c(B_ap \vee B_a\neg p) \wedge B_c(B_aq \vee B_a\neg q)$.

Proposition 7.10. *There is no formula $\varphi \in \mathcal{L}^\div$ such that $\mathcal{M}^\div \varphi \equiv \mathcal{N}$.*

Proof. Suppose towards a contradiction that $\mathcal{M}^\div \varphi \equiv \mathcal{N}$ for some $\varphi \in \mathcal{L}^\div$. Since every world in $\mathcal{M}$ satisfies a distinct set of propositional formulas and contraction preserves propositional valuation, it follows that for all $x \in W$, $(\mathcal{M}^\div \varphi, x) \equiv (\mathcal{N}, x)$. Since $u \notin R_a(w)$ but $u \in R_a^\div \varphi(w)$, we have $\mathcal{M}, w \models B_a \varphi$ and $\mathcal{M}, u \models \neg \varphi$. Similarly, $\mathcal{M}, v \models B_a \varphi$ and $\mathcal{M}, z \models \neg \varphi$. By the definition of contraction, $\{x \in W : \mathcal{M}, x \models \neg \varphi\} \subseteq R_a^\div \varphi(w)$. So $z \in R_a^\div \varphi(w)$. So $\mathcal{M}^\div \varphi, w \models \neg B_a q$. Since $\mathcal{N}, w \models B_a q$, this contradicts the assumption that $(\mathcal{M}^\div \varphi, w)$ and $(\mathcal{N}, w)$ are modally equivalent. $\square$

One way of addressing this problem within the HPAL framework is to restrict the set of worlds that Alice starts considering to a subset of p -possibilities, e.g. those that are compatible with her background knowledge or “entrenched” beliefs. We can also model these dynamics using GDEL. Consider $\mathcal{E} = (E, Q, Q^+, pre)$ depicted as in Figure 9, which captures three aspects of the update: (i) Bob's announcement doesn't eliminate any possibilities (so the Q -relations are universal); (ii) if Alice believes p , then Bob's announcement doesn't make her consider new possibilities ($Q_a^+(e) = Q_a^+(g) = \emptyset$) (iii) if Alice believes $\neg p$, then Bob's announcement makes her consider p possible without revising her beliefs about q ($Q_a^+(f) = \{e\}$ and $Q_a^+(h) = \{g\}$). It's not hard to check that $\mathcal{M} \otimes \mathcal{E} = \mathcal{N}$.




Figure 9: The generalized event model $\mathcal{E} = (E, Q, Q^+, pre)$ representing Bob's announcement to Alice and Clark. As before, we omit the Q -edges, since $Q_i(x) = E$ for all events $x \in E$ and agents $i \in \{a, c\}$. Each event's precondition is given by the conjunction of the literals listed inside the event. So for example event e has precondition $p \wedge q$.

8 Conclusion and future work

In this paper, we introduced a logic for belief contraction due to hedged public announcements, and then a generalized DEL theory for belief expansion and contraction. There are many questions left open for future work, such as: the full characterization of successful formulas for HPAL, the interaction between contraction and expansion operators, properties of a revision operator defined in terms of their combinations via the Levi identity [25], and the iterative properties of GDEL. Of special importance are the closure properties of GDEL. While updates with standard event models and plausibility event models preserve natural properties of accessibility relations, such as transitivity and Euclideaness, this is not true of generalized product updates, even if both the Q - and the Q^+ -relations in the generalized event model are equivalence relations. This has the undesirable consequence that an agent who had full introspection of her own beliefs prior to the update may have higher-order uncertainties about her own beliefs after the update. While it is possible to enforce preservation of introspection in our framework, it is an open question for which class of generalized event models this can be guaranteed.

In addition to the above questions, we also plan to investigate how our model compares with alternative models of belief contraction [28, 24] as well as other models of information loss, such as forgetting and awareness growth [17, 11], and the dynamic approach to interpreting conditionals with modal antecedents. In particular, one may ask how expressive our generalized updates are with respect to plausibility updates. In one sense, we can transform any distinguished Kripke model $\mathcal{K} = (W, R, V)$ into any Kripke model $\mathcal{K}' = (W', R', V')$ with $W = W'$ and $V = V'$ by adding the edges in $\mathcal{K}'$ without keeping any edges in $\mathcal{K}$.¹⁴ To this extent, our framework can emulate the belief dynamics that can be modeled in the plausibility framework, though we'll leave a detailed comparison to future works.

Acknowledgments

We would like to acknowledge Hans van Ditmarsch for suggesting the paper's title and for very helpful comments on the paper, in particular for noticing that the axiomatization of HPAL was missing replacement of equivalents. Gaia would also like to acknowledge Thomas Bolander for very helpful initial discussions on this work and contraction in DEL. We also thank three anonymous reviewers for very helpful comments. Gaia Belardinelli is funded by Independent Research Fund Denmark (grant no. 4255-00020B).

A Proofs

Throughout, by *logical equivalence* we mean provable in **K**. We let $\llbracket \varphi \rrbracket = \{w \in W : \mathcal{M}, w \models \varphi\}$, and $\llbracket \varphi \rrbracket^{\div \Psi} = \{w \in W^{\div \Psi} : \mathcal{M}^{\div \Psi}, w \models \varphi\}$.

Definition A.1 (Bisimulation, refinement, simulation). Let $\mathcal{M} = (W, R, V)$ and $\mathcal{M}' = (W', R', V')$ be two Kripke models. A *bisimulation* between $\mathcal{M}$ and $\mathcal{M}'$ is a non-empty relation $Z \subseteq W \times W'$ such that for all $(w, w') \in Z$ and $a \in \text{Ag}$:

$$(\text{Atom}): w \in V(p) \text{ iff } w' \in V'(p), \text{ for all } p \in \text{At};$$

¹⁴Recall that $\mathcal{K} = (W, R, V)$ is *distinguished* if for every world w there is a unique formula φ such that $\mathcal{M}, w \models \varphi$. The simulation strategy is similar to the strategy used in [21] to show that in DEL with postconditions, one can transform any Kripke model in any other Kripke model. The strategy there is also to get rid of the structure of the initial Kripke model and then reconstruct it by means of carefully devised event models with postconditions.

(Forth): If $v \in R_a(w)$ then there exists $v' \in W'$ such that $v' \in R'_a(w')$ and $(v, v') \in Z$;

(Back): If $v' \in R'_a(w')$ then there exists $v \in W$ such that $v \in R_a(w)$ and $(v, v') \in Z$;

When a bisimulation exists between two models, we say that they are *bisimilar*. A relation that satisfies Atom and Back is a *refinement*. When a refinement exists between $\mathcal{M}$ and $\mathcal{M}'$, we say that $\mathcal{M}'$ is a *refinement* of $\mathcal{M}$. A relation that satisfies Atom and Forth is a *simulation*. When a simulation exists between $\mathcal{M}$ and $\mathcal{M}'$, we say that $\mathcal{M}'$ is a *simulation* of $\mathcal{M}$.

Proof of Theorem 4.2. For soundness, the only non-trivial axiom is A5, namely $[\div \varphi]B_a\psi \leftrightarrow (B_a[\div \varphi]\psi \wedge (B_a\varphi \rightarrow \forall(\neg\varphi \rightarrow [\div \varphi]\psi)))$. Suppose $\mathcal{M}, w \models [\div \varphi]B_a\psi$. We claim that $\mathcal{M}, w \models B_a[\div \varphi]\psi$. Let $v \in R_a(w)$. Then $v \in R_a^{\div \varphi}(w)$. By assumption, $\mathcal{M}^{\div \varphi}, v \models B_a\psi$. So $\mathcal{M}^{\div \varphi}, v \models \psi$. Thus $\mathcal{M}, v \models [\div \varphi]\psi$. So $\mathcal{M}, w \models B_a[\div \varphi]\psi$. Next, we claim that, if $\mathcal{M}, w \models B_a\varphi$ and $\mathcal{M}, w \models [\div \varphi]B_a\psi$, then for every $v \in W$, $\mathcal{M}, v \models \neg\varphi \rightarrow [\div \varphi]\psi$. Let $v \in W$. Suppose $\mathcal{M}, v \models \neg\varphi$. Since $\mathcal{M}, w \models B_a\varphi$, $v \in R^{\div \varphi}(w)$. Since $\mathcal{M}, w \models [\div \varphi]B_a\psi$, $\mathcal{M}^{\div \varphi}, v \models \psi$. So $\mathcal{M}, v \models [\div \varphi]\psi$.

Conversely, suppose $\mathcal{M}, w \models B_a[\div \varphi]\psi \wedge (B_a\varphi \rightarrow \forall(\neg\varphi \rightarrow [\div \varphi]\psi))$. There are two cases, either $\mathcal{M}, w \models \neg B_a\varphi$ or $\mathcal{M}, w \models B_a\varphi$. Suppose $\mathcal{M}, w \models \neg B_a\varphi$. Let $v \in R^{\div \varphi}(w)$. Then $v \in R_a(w)$. Since $\mathcal{M}, w \models B_a[\div \varphi]\psi$, we have $\mathcal{M}, v \models [\div \varphi]\psi$. So $\mathcal{M}^{\div \varphi}, v \models \psi$. Thus $\mathcal{M}, w \models [\div \varphi]B_a\psi$. Now suppose $\mathcal{M}, w \models B_a\varphi$. So $\mathcal{M}, w \models \forall(\neg\varphi \rightarrow [\div \varphi]\psi)$. Let $v \in R^{\div \varphi}(w)$. Either $v \in R_a(w)$ or $\mathcal{M}, v \models \neg\varphi$. In the first case, by the assumption that $\mathcal{M}, w \models B_a[\div \varphi]\psi$, we have $\mathcal{M}^{\div \varphi}, v \models \psi$. In the second case, since $\mathcal{M}, w \models \forall(\neg\varphi \rightarrow [\div \varphi]\psi)$, we have $\mathcal{M}, v \models [\div \varphi]\psi$ or equivalently $\mathcal{M}^{\div \varphi}, v \models \psi$. So $\mathcal{M}, w \models [\div \varphi]B_a\psi$.

We now show that RE is validity preserving. Suppose $\models \varphi \leftrightarrow \psi$. We proceed by induction on the complexity of the environment χ . If χ is a propositional variable $p \in At$, then $\chi[\varphi/p] = \varphi$ and $\chi[\psi/p] = \psi$. So $\models \chi[\varphi/p] \leftrightarrow \chi[\psi/p]$ by assumption. The cases of negation and conjunction follow from propositional logic. Suppose $\models \chi[\varphi/p] \leftrightarrow \chi[\psi/p]$. Then

$$\begin{aligned} \mathcal{M}, w \models B_a\chi[\varphi/p] & \text{ iff } \text{for all } v \in R_a(w) \mathcal{M}, v \models \chi[\varphi/p] \\ & \text{ iff } \text{for all } v \in R_a(w) \mathcal{M}, v \models \chi[\psi/p] \\ & \text{ iff } \mathcal{M}, w \models B_a\chi[\psi/p] \end{aligned}$$

A similar argument shows that $\models \forall\chi[\varphi/p] \leftrightarrow \forall\chi[\psi/p]$. Let $\lambda \in \mathcal{L}^{\div}$. Then

$$\begin{aligned} \mathcal{M}, w \models [\div \lambda]\chi[\varphi/p] & \text{ iff } \mathcal{M}^{\div \lambda}, w \models \chi[\varphi/p] \\ & \text{ iff } \mathcal{M}^{\div \lambda}, w \models \chi[\psi/p] \\ & \text{ iff } \mathcal{M}, w \models [\div \lambda]\chi[\psi/p] \end{aligned}$$

Lastly, for any Kripke model $\mathcal{M} = (W, R, V)$, we claim that $\mathcal{M}^{\div \chi[\varphi/p]} = \mathcal{M}^{\div \chi[\psi/p]}$. It suffices to show that $R_a^{\div \chi[\varphi/p]}(w) = R_a^{\div \chi[\psi/p]}(w)$ for all $w \in W$ and $a \in Ag$. Suppose $\mathcal{M}, w \models \neg B_a\chi[\varphi/p]$. Then by the inductive hypothesis, $\mathcal{M}, w \models \neg B_a\chi[\psi/p]$ and so $R_a^{\div \chi[\varphi/p]}(w) = R_a(w) = R_a^{\div \chi[\psi/p]}(w)$. Suppose $\mathcal{M}, w \models B_a\chi[\varphi/p]$. Then by IH, $\mathcal{M}, w \models B_a\chi[\psi/p]$ and

$$\begin{aligned} R_a^{\div \chi[\varphi/p]}(w) &= R_a(w) \cup \{v \in W : \mathcal{M}, v \models \neg\chi[\varphi/p]\} \\ &= R_a(w) \cup \{v \in W : \mathcal{M}, v \models \neg\chi[\psi/p]\} \\ &= R_a^{\div \chi[\psi/p]}(w) \end{aligned}$$

So $\mathcal{M}^{\div \chi[\varphi/p]} = \mathcal{M}^{\div \chi[\psi/p]}$. Thus, for any Kripke model $\mathcal{M}$ and world w ,

$$\begin{aligned}
\mathcal{M}, w \models [\div \chi[\varphi/p]]\lambda & \text{ iff } \mathcal{M}^{\div \chi[\varphi/p]}, w \models \lambda \\
& \text{ iff } \mathcal{M}^{\div \chi[\psi/p]}, w \models \lambda \\
& \text{ iff } \mathcal{M}, w \models [\div \chi[\psi/p]]\lambda
\end{aligned}$$

Given the reduction axioms, every formula with dynamic modalities is provably equivalent in HPAL to a formula without dynamic modalities. So the completeness of HPAL follows from the completeness of the basic doxastic logic with universal modalities. $\square$

Proof of Theorem 5.5. Suppose $Ag = \{a\}$. Let $\varphi \in \mathcal{L}$. Let $\mathcal{M} = (W, R, V)$ be a Kripke model that validates **S5**. Suppose $\mathcal{M}, w \models \varphi$. Since R_a is an equivalence relation, for all $v \in R_a(w)$, $\mathcal{M}, v \models \tilde{B}_a \varphi$. Then $R_a(v) = R_a^{\div \neg \varphi}(v)$ for all $v \in R_a(w)$. Consider $Z = \{(v, v) : v \in R_a(w)\}$. We claim that Z is a bisimulation between $(\mathcal{M}, w)$ and $(\mathcal{M}^{\div \neg \varphi}, w)$. Clearly the two worlds satisfy the same proposition atoms. Let $(v, v) \in Z$. Then $v \in R_a(w)$. Suppose $u \in R_a(v)$. Since R_a is an equivalence relation, we have $u \in R_a(w)$ and so $(u, u) \in Z$. Moreover, since $R_a(v) = R_a(w) = R_a^{\div \neg \varphi}(w) = R_a^{\div \neg \varphi}(v)$, we have $u \in R_a^{\div \neg \varphi}(v)$. Similarly, suppose $u \in R_a^{\div \neg \varphi}(v)$. Then $u \in R_a^{\div \neg \varphi}(w) = R_a(w) = R_a(v)$. So Z is a bisimulation between $(\mathcal{M}, w)$ and $(\mathcal{M}^{\div \neg \varphi}, w)$. Thus, $\mathcal{M}^{\div \neg \varphi}, w \models \varphi$. $\square$

Proof of Proposition 6.3. Let $\varphi, \psi, \chi, \lambda \in \mathcal{L}^{\div}$ with χ, λ propositional and consider a Kripke model $\mathcal{M} = (W, R, V)$ and a world w in it. Notationally, let $\mathcal{M}^{\div \varphi} = (W^{\div \varphi}, R^{\div \varphi}, V^{\div \varphi})$, $\llbracket \psi \rrbracket = \{w \in W : \mathcal{M}, w \models \psi\}$ and $\llbracket \psi \rrbracket^{\div \varphi} = \{w \in W : \mathcal{M}^{\div \varphi}, w \models \psi\}$.

1. *Propositional Inclusion:* Assume that $\mathcal{M}, w \models [\div \varphi]B_a \chi$. Then $\mathcal{M}^{\div \varphi}, w \models B_a \chi$. So $R_a^{\div \varphi}(w) \subseteq \llbracket \chi \rrbracket^{\div \varphi}$. By definition of contraction $R_a(w) \subseteq R_a^{\div \varphi}(w) \subseteq \llbracket \chi \rrbracket^{\div \varphi}$. As χ is propositional, $\llbracket \chi \rrbracket = \llbracket \chi \rrbracket^{\div \varphi}$. Hence, $\mathcal{M}, w \models B_a \chi$.
2. *Propositional Conjunctive Inclusion:* Suppose $\mathcal{M}, w \models [\div \lambda \wedge \psi] \neg B_a \lambda$ and $\mathcal{M}, w \models [\div \lambda \wedge \psi] B_a \chi$. There are two cases: either $\mathcal{M}, w \models \neg B_a(\lambda \wedge \psi)$ or $\mathcal{M}, w \models B_a(\lambda \wedge \psi)$. Suppose the first. Then $R_a^{\div \lambda \wedge \psi}(w) = R_a(w)$. Since $\mathcal{M}, w \models [\div \lambda \wedge \psi] \neg B_a \lambda$, there exists $v \in R_a^{\div \lambda \wedge \psi}(w)$ with $\mathcal{M}^{\div \lambda \wedge \psi}, v \models \neg \lambda$. As λ is propositional, then $\mathcal{M}, v \models \neg \lambda$, and since $R_a^{\div \lambda \wedge \psi}(w) = R_a(w)$ then $\mathcal{M}, w \models \neg B_a \lambda$. So $R_a^{\div \lambda}(w) = R_a(w)$. Hence, $R_a^{\div \lambda \wedge \psi}(w) = R_a^{\div \lambda}(w)$. So $\mathcal{M}, w \models [\div \lambda \wedge \psi] B_a \chi$ implies $R_a^{\div \lambda}(w) = R_a^{\div \lambda \wedge \psi}(w) \subseteq \llbracket \chi \rrbracket^{\div \lambda \wedge \psi} = \llbracket \chi \rrbracket^{\div \lambda}$, where the last identity holds as χ is propositional. Hence, $\mathcal{M}, w \models [\div \lambda] B_a \chi$. Now suppose the second case holds. Then $\mathcal{M}, w \models B_a \lambda$, and so $R_a^{\div \lambda}(w) = R_a(w) \cup \llbracket \neg \lambda \rrbracket \subseteq R_a(w) \cup \llbracket \neg \lambda \vee \neg \psi \rrbracket = R_a^{\div (\lambda \wedge \psi)}(w)$. As by assumption $\mathcal{M}, w \models [\div \lambda \wedge \psi] B_a \chi$, then $R_a^{\div (\lambda \wedge \psi)}(w) \subseteq \llbracket \chi \rrbracket^{\div (\lambda \wedge \psi)} = \llbracket \chi \rrbracket^{\div \lambda}$. Hence, $R_a^{\div \lambda}(w) \subseteq \llbracket \chi \rrbracket^{\div \lambda}$ and so $\mathcal{M}, w \models [\div \lambda] B_a \chi$.
3. *Propositional Conjunctive Overlap.* Suppose that $\mathcal{M}, w \models [\div \varphi] B_a \chi \wedge [\div \psi] B_a \chi$. Then $R_a^{\div \varphi}(w) \subseteq \llbracket \chi \rrbracket^{\div \varphi}$ and $R_a^{\div \psi}(w) \subseteq \llbracket \chi \rrbracket^{\div \psi}$. As χ is propositional, $\llbracket \chi \rrbracket^{\div \varphi} = \llbracket \chi \rrbracket^{\div \psi} = \llbracket \chi \rrbracket^{\div \varphi \wedge \psi}$, and so $(R_a^{\div \varphi}(w) \cup R_a^{\div \psi}(w)) \subseteq \llbracket \chi \rrbracket^{\div \varphi \wedge \psi}$. Two cases: either $\mathcal{M}, w \models B_a(\varphi \wedge \psi)$ or not. If the first, then $\mathcal{M}, w \models B_a \varphi$ and $\mathcal{M}, w \models B_a \psi$. So $R_a^{\div \varphi}(w) = R_a(w) \cup \llbracket \neg \varphi \rrbracket$ and $R_a^{\div \psi}(w) = R_a(w) \cup \llbracket \neg \psi \rrbracket$ and $R_a^{\div \varphi \wedge \psi}(w) = R_a(w) \cup \llbracket \neg \varphi \vee \neg \psi \rrbracket = R_a(w) \cup \llbracket \neg \varphi \rrbracket \cup \llbracket \neg \psi \rrbracket$. Then, $R_a^{\div \varphi \wedge \psi}(w) = R_a^{\div \varphi}(w) \cup R_a^{\div \psi}(w)$, and so $R_a^{\div \varphi \wedge \psi}(w) \subseteq \llbracket \chi \rrbracket^{\div \varphi \wedge \psi}$. Hence, $\mathcal{M}, w \models [\div \varphi \wedge \psi] B_a \chi$. Suppose instead that $\mathcal{M}, w \models \neg B_a(\varphi \wedge \psi)$. Then, $R_a^{\div \varphi \wedge \psi}(w) = R_a(w)$, and since by initial assumption $R_a^{\div \varphi}(w) \subseteq \llbracket \chi \rrbracket^{\div \varphi \wedge \psi}$ and by definition of contraction $R_a(w) \subseteq R_a^{\div \varphi}(w)$, then $R_a(w) \subseteq \llbracket \chi \rrbracket^{\div \varphi \wedge \psi}$. Hence, $\mathcal{M}, w \models [\div \varphi \wedge \psi] B_a \chi$. $\square$

Proof of Proposition 7.4. Let $\mathcal{M} = (W, R, V)$ be a Kripke model. Notationally, we let $\mathcal{E}(\div \varphi) = (E, Q, Q^+, pre)$, $\mathcal{M} \otimes \mathcal{E}(\div \varphi) = (W^{\mathcal{E}(\div \varphi)}, R^{\mathcal{E}(\div \varphi)}, V^{\mathcal{E}(\div \varphi)})$ and $\mathcal{M}^{\div \varphi} = (W^{\div \varphi}, R^{\div \varphi}, V^{\div \varphi})$. The function $h: W^{\mathcal{E}(\div \varphi)} \rightarrow W^{\div \varphi}$, defined by $h((w, e)) = w$, for all $w \in W$, is a bijection: first, notice that, by

definition of $\mathcal{E}(\div \varphi)$, for each world $w \in W$ there is exactly one event $e \in E$ such that $\mathcal{M}, w \models \text{pre}(e)$, and so $(w, e) \in W^{\mathcal{E}(\div \varphi)}$ iff $w \in W$. Then notice that, by definition of $\div \varphi$ -update, $W^{\div \varphi} = W$, and so $(w, e) \in W^{\mathcal{E}(\div \varphi)}$ iff $w \in W^{\div \varphi}$. We now show that h preserves atomic valuations and accessibility relations. The atomic valuations are clearly preserved, as $V^{\mathcal{E}(\div \varphi)}(p) = V(p) = V^{\div \varphi}(p)$ for all $p \in \text{At}$. For the accessibility relations, let $a \in \text{Ag}$ and $(w, e) \in W^{\mathcal{E}(\div \varphi)}$. We want to show that $(v, f) \in R_a^{\mathcal{E}(\div \varphi)}(w, e)$ iff $v \in R_a^{\div \varphi}(w)$. We consider two cases: either $\mathcal{M}, w \not\models B_a \varphi$, or $\mathcal{M}, w \models B_a \varphi$. Suppose $\mathcal{M}, w \not\models B_a \varphi$. Then $\mathcal{E}(\div \varphi)$ is such that $Q_a(e) = E$ and $Q_a^+(e) = \emptyset$, and $\mathcal{M}^{\div \varphi}$ is such that $R_a(w) = R_a^{\div \varphi}(w)$. So we have that $(v, f) \in R_a^{\mathcal{E}(\div \varphi)}(w, e)$ iff (by definition of generalized product update) $v \in R_a(w)$ iff (by definition of $\div \varphi$ -update) $v \in R_a^{\div \varphi}(w)$, as we wanted to show. Suppose instead $\mathcal{M}, w \models B_a \varphi$. Assume $(v, f) \in R_a^{\mathcal{E}(\div \varphi)}(w, e)$. Then either $v \in R_a(w)$ and $f \in Q_a(e)$, or $f \in Q_a^+(e)$. If $v \in R_a(w)$, then $v \in R_a^{\div \varphi}(w)$, as desired. If instead $f \in Q_a^+(e)$, then $\mathcal{M}, v \models \neg \varphi$ and so $v \in R_a^{\div \varphi}(w)$, as desired. For the other direction, assume $v \in R_a^{\div \varphi}(w)$. Then either $v \in R_a(w)$ or $\mathcal{M}, v \models \neg \varphi$. If $v \in R_a(w)$, then since $Q_a(e) = E$, $(v, f) \in R_a^{\mathcal{E}(\div \varphi)}(w, e)$. If instead $\mathcal{M}, v \models \neg \varphi$, then by $(v, f) \in W^{\mathcal{E}(\div \varphi)}$, $\text{pre}(f) \models \neg \varphi$ by definition of $\mathcal{E}(\div \varphi)$. Since $\mathcal{M}, w \models B_a \varphi$ then $Q_a^+(e) = \{g \in E : \text{pre}(g) \models \neg \varphi\}$, and so $(v, f) \in R_a^{\mathcal{E}(\div \varphi)}(w, e)$. Hence in all cases we have the desired. $\square$

Proof of Theorem 7.7. It is straightforward that the inference rules preserve validity. We check the soundness of A4 and A5. The soundness of the other axioms is immediate. Let $\mathcal{M} = (W, R, V)$ be a Kripke model, let $\mathcal{E} = (E, Q, Q^+, \text{pre})$ be a generalized event model, and let $\mathcal{M} \otimes \mathcal{E} = (W^{\mathcal{E}}, R^{\mathcal{E}}, V^{\mathcal{E}})$.

- A4: Note that, if $\mathcal{M}, w \not\models \text{pre}(e)$, then the biconditional holds automatically at $(\mathcal{M}, w)$. Suppose $\mathcal{M}, w \models \text{pre}(e)$ and $\mathcal{M}, w \models [\mathcal{E}, e] \forall \varphi$. Then $\mathcal{M} \otimes \mathcal{E}, (w, e) \models \forall \varphi$, i.e. for all $(v, f) \in W^{\mathcal{E}}$, $\mathcal{M} \otimes \mathcal{E}, (v, f) \models \varphi$. In other words, for all $f \in E$ and $v \in W$ such that $\mathcal{M}, v \models \text{pre}(f)$, $\mathcal{M}, v \models [\mathcal{E}, f] \varphi$. So $\mathcal{M}, w \models \bigwedge_{f \in E} \forall [\mathcal{E}, f] \varphi$.

Conversely, suppose $\mathcal{M}, w \models \text{pre}(e)$ but $\mathcal{M}, w \not\models \neg [\mathcal{E}, e] \forall \varphi$, i.e. $\mathcal{M} \otimes \mathcal{E}, (w, e) \not\models \forall \varphi$. So there is $(v, f) \in W^{\mathcal{E}}$ such that $\mathcal{M} \otimes \mathcal{E}, (v, f) \not\models \varphi$. It follows that $\mathcal{M}, v \models \neg [\mathcal{E}, f] \varphi$ and so $\mathcal{M}, w \models \neg \bigwedge_{f \in E} \forall [\mathcal{E}, f] \varphi$.

- A5: Suppose $\mathcal{M}, w \models [\mathcal{E}, e] B_a \varphi \wedge \text{pre}(e)$. Let $f \in Q_a^+(e)$. Then for any $v \in W$, if $\mathcal{M}, v \models \text{pre}(f)$, we have $(v, f) \in R_a^{\mathcal{E}}((w, e))$ and so by assumption $\mathcal{M}^{\mathcal{E}}, (v, f) \models \varphi$. Thus $\mathcal{M}, w \models \bigwedge_{f \in Q_a^+(e)} \forall [\mathcal{E}, f] \varphi$. Let $f \in Q_a(e)$ and $v \in R_a(w)$. Then $(v, f) \in R_a^{\mathcal{E}}((w, e))$ and so $\mathcal{M}^{\mathcal{E}}, (v, f) \models \varphi$. Thus $\mathcal{M}, v \models [\mathcal{E}, f] \varphi$. So $\mathcal{M}, w \models \bigwedge_{f \in Q_a(e)} B_a [\mathcal{E}, f] \varphi$.

Conversely, suppose $\mathcal{M}, w \models \text{pre}(e) \rightarrow \bigwedge_{f \in Q_a^+(e)} \forall [\mathcal{E}, f] \varphi \wedge \bigwedge_{f \in Q_a(e)} B_a [\mathcal{E}, f] \varphi$ and let $(v, f) \in R_a^{\mathcal{E}}((w, e))$. Either $f \in Q_a^+(e)$ or $f \in Q_a(e)$ and $v \in R_a(w)$. In the first case, since $\mathcal{M}, w \models \forall [\mathcal{E}, f] \varphi$, we have $\mathcal{M}, v \models [\mathcal{E}, f] \varphi$. In the second case, since $\mathcal{M}, w \models B_a [\mathcal{E}, f] \varphi$, we have $\mathcal{M}, v \models [\mathcal{E}, f] \varphi$. So either way, $\mathcal{M}^{\mathcal{E}}, (v, f) \models \varphi$ and so $\mathcal{M}^{\mathcal{E}}, (w, e) \models B_a \varphi$.

- A6: The proof mainly requires showing that there exists an isomorphism between $(\mathcal{M} \otimes \mathcal{E}) \otimes \mathcal{F} = (W^{\mathcal{E} \circ \mathcal{F}}, R^{\mathcal{E} \circ \mathcal{F}}, V^{\mathcal{E} \circ \mathcal{F}})$ and $\mathcal{M} \otimes (\mathcal{E} \circ \mathcal{F}) = (W^{\mathcal{E} \circ \mathcal{F}}, R^{\mathcal{E} \circ \mathcal{F}}, V^{\mathcal{E} \circ \mathcal{F}})$. By invariance of truth under isomorphism, this will imply the desired. We use notation $\mathcal{M} \otimes \mathcal{E} = (W^{\mathcal{E}}, R^{\mathcal{E}}, V^{\mathcal{E}})$, and $\mathcal{E} = (E^{\mathcal{E}}, Q^{\mathcal{E}}, Q^{+\mathcal{E}}, \text{pre}^{\mathcal{E}})$ and $\mathcal{F} = (E^{\mathcal{F}}, Q^{\mathcal{F}}, Q^{+\mathcal{F}}, \text{pre}^{\mathcal{F}})$.

To show the existence of the isomorphism, let $h : W^{\mathcal{E} \circ \mathcal{F}} \rightarrow W^{\mathcal{E} \circ \mathcal{F}}$ be defined by $h((w, e), f) = (w, (e, f))$. Clearly, h is a bijection: $((w, e), f) \in W^{\mathcal{E} \circ \mathcal{F}}$ iff $\mathcal{M}, w \models \text{pre}^{\mathcal{E}}(e)$ and $\mathcal{M} \otimes \mathcal{E}, (w, e) \models \text{pre}^{\mathcal{F}}(f)$, iff $\mathcal{M}, w \models \text{pre}^{\mathcal{E}}(e) \wedge [\mathcal{E}, e] \text{pre}^{\mathcal{F}}(f)$, iff $(w, (e, f)) \in W^{\mathcal{E} \circ \mathcal{F}}$. It also preserves valuations, since for every atom $p \in \text{At}$, $((w, e), f) \in V^{\mathcal{E} \circ \mathcal{F}}(p)$ iff $w \in V(p)$ iff $(w, (e, f)) \in V^{\mathcal{E} \circ \mathcal{F}}(p)$. Finally, h preserves accessibility relations. Let $((w', e'), f') \in R_a^{\mathcal{E} \circ \mathcal{F}}((w, e), f)$. We want to show that $(w', (e', f')) \in R_a^{\mathcal{E} \circ \mathcal{F}}((w, (e, f)))$. Let $((w', e'), f') \in R_a^{\mathcal{E} \circ \mathcal{F}}((w, (e, f)))$. Two cases: either $f' \in$

$Q_a^{+\mathcal{F}}(f)$ or $(w', e') \in R_a^\mathcal{E}((w, e))$ and $f' \in Q_a^\mathcal{F}(f)$. If $f' \in Q_a^{+\mathcal{F}}(f)$ then $(e', f') \in Q_a^{+\mathcal{E} \circ \mathcal{F}}(e, f)$ and so $(w', (e', f')) \in R_a^{\mathcal{E} \circ \mathcal{F}}((w, (e, f)))$. If $(w', e') \in R_a^\mathcal{E}((w, e))$ and $f' \in Q_a^\mathcal{F}(f)$, then again two cases, either $e' \in Q_a^{+\mathcal{E}}(e)$, or $w' \in R_a(w)$ and $e' \in Q_a^\mathcal{E}(e)$. If $e' \in Q_a^{+\mathcal{E}}(e)$ then by $f' \in Q_a^\mathcal{F}(f)$ and definition of composition, $(e', f') \in Q_a^{+\mathcal{E} \circ \mathcal{F}}((e, f))$, which implies $(w', (e', f')) \in R_a^{\mathcal{E} \circ \mathcal{F}}((w, (e, f)))$. If $w' \in R_a(w)$ and $e' \in Q_a^\mathcal{E}(e)$, then by $f' \in Q_a^\mathcal{F}(f)$ and definition of composition we have $(e', f') \in Q_a^{\mathcal{E} \circ \mathcal{F}}((e, f))$, and so $(w', (e', f')) \in R_a^{\mathcal{E} \circ \mathcal{F}}((w, (e, f)))$. Hence, in all cases the desired holds. The other direction (that is, if $(w', (e', f')) \in R_a^{\mathcal{E} \circ \mathcal{F}}((w, (e, f)))$ then $((w', e'), f') \in R_a^\mathcal{F}((w, e), f)$) follows by a similar straightforward application of the definition of composition and generalized product update.

Thus, h is an isomorphism. By invariance of truth under isomorphisms, for every $((w, e), f) \in W^{\mathcal{E} \circ \mathcal{F}}$ and every $\varphi \in \mathcal{L}_{\text{DEL}}$, we have $((\mathcal{M} \otimes \mathcal{E}) \otimes \mathcal{F}, ((w, e), f)) \models \varphi$ iff $(\mathcal{M} \otimes (\mathcal{E} \circ \mathcal{F}), (w, (e, f))) \models \varphi$. It remains only to connect this with the truth conditions for the dynamic modalities. Let $w \in W$. If $\mathcal{M}, w \not\models \text{pre}^\mathcal{E}(e)$, then both $[\mathcal{E}, e][\mathcal{F}, f]\varphi$ and $[\mathcal{E} \circ \mathcal{F}, (e, f)]\varphi$ are true at w vacuously, since $\text{pre}^{\mathcal{E} \circ \mathcal{F}}(e, f) = \text{pre}^\mathcal{E}(e) \wedge [\mathcal{E}, e]\text{pre}^\mathcal{F}(f)$. If $\mathcal{M}, w \models \text{pre}^\mathcal{E}(e)$ but $\mathcal{M} \otimes \mathcal{E}, (w, e) \not\models \text{pre}^\mathcal{F}(f)$, then again $[\mathcal{E}, e][\mathcal{F}, f]\varphi$ is true at w vacuously, and so is $[\mathcal{E} \circ \mathcal{F}, (e, f)]\varphi$, since $\mathcal{M}, w \not\models \text{pre}^{\mathcal{E} \circ \mathcal{F}}(e, f)$. Finally, if $\mathcal{M}, w \models \text{pre}^\mathcal{E}(e)$ and $\mathcal{M} \otimes \mathcal{E}, (w, e) \models \text{pre}^\mathcal{F}(f)$, then both product points $((w, e), f)$ and $(w, (e, f))$ exist, and the desired equivalence follows from the isomorphism above. Hence, in all cases, $\mathcal{M}, w \models [\mathcal{E}, e][\mathcal{F}, f]\varphi$ iff $\mathcal{M}, w \models [\mathcal{E} \circ \mathcal{F}, (e, f)]\varphi$. Since $\mathcal{M}$ and w were arbitrary, the equivalence is valid.

Completeness of GDEL follows from the completeness of K with the universal modality by standard reduction arguments. $\square$

Proof of Theorem 7.8. Let $\mathcal{M} = (W, R, V)$ be a Kripke model and $\mathcal{E} = (E, Q, Q^+, \text{pre})$ be a generalized event model. Notationally, let $\mathcal{M} \otimes \mathcal{E} = (W^\mathcal{E}, R^\mathcal{E}, V^\mathcal{E})$. Our proof strategy is to define a submodel of $\mathcal{M} \otimes \mathcal{E}$ that serves as the intermediate refinement step, and then show that $\mathcal{M} \otimes \mathcal{E}$ is a simulation of that model. Let $(W^{\mathcal{E}^-}, R^{\mathcal{E}^-}, V^{\mathcal{E}^-})$ be a Kripke model such that $W^{\mathcal{E}^-} = W^\mathcal{E}$, $R_a^{\mathcal{E}^-}(w, e) = \{(v, f) \in W^\mathcal{E} : v \in R_a(w), f \in Q_a(e)\}$, for all $a \in \text{Ag}$, and $V^{\mathcal{E}^-} = V^\mathcal{E}$. This is a submodel of $\mathcal{M} \otimes \mathcal{E}$, only missing the edges between worlds that the Q^+ relation added. It is then clear that we recover the full $R^\mathcal{E}$ by taking the union of $R^{\mathcal{E}^-}$ with the set containing those edges, that is $R_a^\mathcal{E}(w, e) = R_a^{\mathcal{E}^-}(w, e) \cup \{(v, f) \in W^\mathcal{E} : f \in Q_a^+(e)\}$, for all $a \in \text{Ag}$ and $(w, e) \in W^\mathcal{E}$. Now, $(W^{\mathcal{E}^-}, R^{\mathcal{E}^-}, V^{\mathcal{E}^-})$ is a standard DEL update, and it is a well known result that it is a refinement of $\mathcal{M}$ [15]. We now show that $\mathcal{M} \otimes \mathcal{E}$ is a simulation of $\mathcal{M} \otimes \mathcal{E}^-$, from which we can conclude that $\mathcal{M} \otimes \mathcal{E}$ is a simulation of a refinement of $\mathcal{M}$. Let $Z \subseteq W^{\mathcal{E}^-} \times W^\mathcal{E}$ be such that $((w, e), (w, e)) \in Z$. As $R^\mathcal{E}$ only adds edges to $R^{\mathcal{E}^-}$ without removing any, it clearly holds that if $(v, f) \in R_a^{\mathcal{E}^-}(w, e)$, then $(v, f) \in R_a^\mathcal{E}(w, e)$, and since $((v, f), (v, f)) \in Z$ then Z is a simulation. $\square$

References

- [1] Carlos E. Alchourrón, Peter Gärdenfors & David Makinson (1985): *On the Logic of Theory Change: Partial Meet Contraction and Revision Functions*. *The Journal of Symbolic Logic* 50(2), pp. 510–530, doi:10.2307/2274239.
- [2] Mikkel Birkegaard Andersen, Thomas Bolander, Hans van Ditmarsch & Martin Holm Jensen (2013): *Bisimulation for single-agent plausibility models*. In: *Australasian Joint Conference on Artificial Intelligence*, Springer, pp. 277–288, doi:10.1007/978-3-319-03680-9_30.

- [3] Mikkel Birkegaard Andersen, Thomas Bolander, Hans van Ditmarsch & Martin Holm Jensen (2017): *Bisimulation and expressivity for conditional belief, degrees of belief, and safe belief*. *Synthese* 194(7), pp. 2447–2487, doi:10.1007/s11229-016-1060-x.
- [4] Hajnal Andr  ka, Istv  n N  meti & Johan Van Benthem (1998): *Modal languages and bounded fragments of predicate logic*. *Journal of philosophical logic* 27(3), pp. 217–274, doi:10.1023/A:1004275029985.
- [5] Guillaume Aucher, Philippe Balbiani, Luis Fari  as del Cerro & Andreas Herzig (2009): *Global and Local Graph Modifiers*. *Electronic Notes in Theoretical Computer Science* 231, pp. 293–307, doi:10.1016/j.entcs.2009.02.042. Proceedings of the 5th Workshop on Methods for Modalities (M4M5 2007).
- [6] Alexandru Baltag, Lawrence S. Moss & Slawomir Solecki (1998): *The Logic of Public Announcements and Common Knowledge and Private Suspicions*. In Itzhak Gilboa, editor: *Proceedings of the 7th Conference on Theoretical Aspects of Rationality and Knowledge (TARK-98)*, Morgan Kaufmann, pp. 43–56.
- [7] Alexandru Baltag & Bryan Renne (2016): *Dynamic Epistemic Logic*. In Edward N. Zalta, editor: *The Stanford Encyclopedia of Philosophy*, Winter 2016 edition, Metaphysics Research Lab, Stanford University. Available at <https://plato.stanford.edu/archives/win2016/entries/dynamic-epistemic/>.
- [8] Alexandru Baltag & Sonja Smets (2008): *A Qualitative Theory of Dynamic Interactive Belief Revision*. In Giacomo Bonanno, Wiebe van der Hoek & Michael Wooldridge, editors: *Logic and the Foundations of Game and Decision Theory (LOFT7)*, *Texts in Logic and Games* 3, Amsterdam University Press, pp. 13–60. Available at <https://www.jstor.org/stable/j.ctt46mz4h.4>.
- [9] Johan van Benthem (2007): *Dynamic logic for belief revision*. *Journal of applied non-classical logics* 17(2), pp. 129–155, doi:10.3166/jancl.17.129-155.
- [10] Johan van Benthem (2023): *The Logic of Conditionals on Outback Trails*. *Logic Journal of the IGPL* 31(6), pp. 1135–1152, doi:10.1093/jigpal/jzac064.
- [11] Johan van Benthem & Fernando R. Vel  zquez-Quesada (2010): *The dynamics of awareness*. *Synthese* 177, pp. 5–27, doi:10.1007/s11229-010-9764-9.
- [12] Johan van Benthem and (2007): *Dynamic logic for belief revision*. *Journal of Applied Non-Classical Logics* 17(2), pp. 129–155, doi:10.3166/jancl.17.129-155.
- [13] Matthew A. Benton & Peter Van Elswyk (2020): *Hedged Assertion*. In Sanford Goldberg, editor: *The Oxford Handbook of Assertion*, Oxford University Press, pp. 245–263, doi:10.1093/oxfordhb/9780190675233.013.11.
- [14] Patrick Blackburn, Maarten de Rijke & Yde Venema (2001): *Modal Logic*. *Cambridge Tracts in Theoretical Computer Science* 53, Cambridge University Press, Cambridge, UK, doi:10.1017/CBO9781107050884.
- [15] Laura Bozzelli, Hans van Ditmarsch, Tim French, James Hales & Sophie Pinchinat (2014): *Refinement modal logic*. *Information and Computation* 239, pp. 303–339, doi:10.1016/j.ic.2014.07.013.
- [16] Lorenz Demey (2011): *Some remarks on the model theory of epistemic plausibility models*. *Journal of Applied Non-Classical Logics* 21(3-4), pp. 375–395, doi:10.3166/jancl.21.375-395.

- [17] Hans van Ditmarsch & Tim French (2009): *Awareness and forgetting of facts and agents*. In: *2009 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology*, 3, IEEE, pp. 478–483, doi:10.1109/WI-IAT.2009.330.
- [18] Hans van Ditmarsch, Tim French, Rustam Galimullin & Louwe B. Kuijer (2025): *Modal Logic for Simulation, Refinement, and Mutual Ignorance*. *Electronic Proceedings in Theoretical Computer Science* 437, p. 379–398, doi:10.4204/eptcs.437.30.
- [19] Hans van Ditmarsch, Andreas Herzig, Jérôme Lang & Pierre Marquis (2009): *Introspective forgetting*. *Synthese* 169(2), pp. 405–423, doi:10.1007/s11229-009-9554-4.
- [20] Hans van Ditmarsch, Wiebe van der Hoek & Barteld Kooi (2008): *Dynamic epistemic logic*. Springer, doi:10.1007/978-1-4020-5839-4.
- [21] Hans van Ditmarsch & Barteld Kooi (2008): *Semantic Results for Ontic and Epistemic Change*. In Giacomo Bonanno, Wiebe van der Hoek & Michael Wooldridge, editors: *Logic and the Foundation of Game and Decision Theory (LOFT 7)*, pp. 87–117. Available at <https://www.jstor.org/stable/j.ctt46mz4h.6>.
- [22] David Fernández-Duque, Ángel Nepomuceno-Fernández, Enrique Sarrión-Morrillo, Fernando Soler-Toscano & Fernando R. Velázquez-Quesada (2015): *Forgetting complex propositions*. *Logic Journal of the IGPL* 23(6), pp. 942–965, doi:10.1093/jigpal/jzv049.
- [23] Virginie Fiutek (2013): *Playing with Knowledge and Belief*. Ph.D. thesis, University of Amsterdam, Institute for Logic, Language, and Computation. Available at <https://eprints.illc.uva.nl/id/eprint/2120>.
- [24] Patrick Girard, Jeremy Seligman & Fenrong Liu (2012): *General Dynamic Dynamic Logic*. In Thomas Bolander, Torben Braüner, Silvio Ghilardi & Lawrence Moss, editors: *Advances in Modal Logic* 9, College Publications, pp. 239–260, doi:10.1093/oso/9780198538592.003.0004.
- [25] Sven Ove Hansson (2022): *Logic of Belief Revision*. In Edward N. Zalta, editor: *The Stanford Encyclopedia of Philosophy*, Spring 2022 edition, Metaphysics Research Lab, Stanford University. Available at <https://plato.stanford.edu/archives/spr2022/entries/logic-belief-revision/>.
- [26] Andreas Herzig (2017): *Dynamic Epistemic Logics: Promises, Problems, Shortcomings, and Perspectives*. *Journal of Applied Non-Classical Logics* 27(3-4), pp. 328–341, doi:10.1080/11663081.2017.1416036.
- [27] Wesley H. Holliday & Thomas F. Icard III (2017): *Indicative conditionals and dynamic epistemic logic*. *arXiv preprint arXiv:1707.08752*, doi:10.48550/arXiv.1707.08752.
- [28] Hannes Leitgeb & Krister Segerberg (2007): *Dynamic doxastic logic: why, how, and where to?* *Synthese* 155(2), pp. 167–190, doi:10.1007/s11229-006-9143-8.
- [29] Jan Plaza (2007): *Logics of public communications*. *Synthese* 158(2), pp. 165–179, doi:10.1007/s11229-007-9168-7.
- [30] Krister Segerberg (1999): *Two Traditions in the Logic of Belief: Bringing them Together*, pp. 135–147. Springer Netherlands, Dordrecht, doi:10.1007/978-94-011-4574-9_8.
- [31] Yanjing Wang & Qinxian Cao (2013): *On axiomatizations of public announcement logic*. *Synthese* 190(Suppl 1), pp. 103–134, doi:10.1007/s11229-012-0233-5.
- [32] Seth Yalcin (2007): *Epistemic modals*. *Mind* 116(464), pp. 983–1026, doi:10.1093/mind/fzm983.