Suspending Judgement: Belief Contraction in Dynamic Epistemic Logic

Alexandru Baltag[1] , **Virginie Fiutek**[2] , **Sonja Smets**[1]
[1]Institute for Logic, Language and Computation; University of Amsterdam
[2]Independent Researcher
TheAlexandruBaltag@gmail.com, fiutek.virginie@gmail.com, S.J.L.Smets@uva.nl

Abstract

We look at multi-agent versions of three different belief contraction operations (severe withdrawal, conservative contraction and moderate contraction), considering them as dynamic operations on plausibility models, meant to represent joint actions of "suspension of belief" by groups of agents (or by individual agents). We provide sound and complete axiomatizations, in the presence of standard static operators such as conditional belief, knowledge and safe belief.

1 Introduction

Belief contraction is meant to be a formalization of the act of "suspending judgment" over a fact ψ: even if ψ was previously believed, the action of contracting with ψ "raises doubt" over its veracity. In this paper we study *multi-agent* contraction in the framework of *'soft' Dynamic Epistemic Logic* (DEL), based on epistemic plausibility models, cf. (Baltag and Smets 2006c; van Benthem 2007; Baltag and Smets 2006b; Baltag and Smets 2008). However, there are significant differences between the traditional framework of Belief Revision Theory and the DEL approach (Baltag, Moss, and Solecki 1998; van Ditmarsch, van der Hoek, and Kooi 2008; Baltag and Moss 2004; van Benthem 2011).

The standard AGM theory for belief revision (Alchourron, Gardenfors, and Makinson 1985) focuses only on the revision of 'first-level' (factual) beliefs, and does not handle the perspective of agents having higher-order 'beliefs about beliefs' (be they their own or others' beliefs). This means that the AGM mechanism for belief revision cannot straightforwardly be cast into an epistemic-doxastic logic for multi-agent systems in which the agents have the capability of higher-order reasoning. For this reason, DEL introduces a distinction between 'static' revision (usually encoded in DEL as conditional beliefs $B^{\psi}\varphi$, see e.g. (Baltag and Smets 2008)) and 'dynamic' revision operators ("doxastic upgrades"). A conditional belief operator is a binary modal operator, written as $B_a^{\psi}\varphi$ to express that agent a believes φ conditional on ψ. It captures agent a's beliefs after revising with ψ about the state of the world before the revision. Indeed, the sentence $B_a^{\psi}\varphi$ states that, if agent a would learn ψ, then she would come to believe that φ was the case (before the learning). In contrast, DEL encodes dynamic belief revision using PDL-style operators such as $[\Uparrow \psi]\varphi$ and $[\uparrow \psi]\varphi$ (for radical and conservative upgrade with ψ, which we will denote here as $[*^r\psi]\varphi$ and $[*^c\psi]\varphi$), that can capture the agents' revised beliefs about the world *as it will be after the revision took place*.[1]

While various dynamic revision operators have been thoroughly studied in DEL (see (Baltag and Smets 2006c; van Benthem 2007; Baltag and Smets 2008)), there seems to exist a still widespread misconception (even among some DEL researchers), according to which belief contraction cannot be easily accommodated in this framework. However, there are a number of successful attempts in this sense, cf. the PhD thesis of (Fiutek 2013) and our related joint work (Baltag, Fiutek, and Smets 2014), as well as the PhD thesis of (Souza 2016) and the follow-up papers (Souza et al. 2016; Souza, Vieira, and Moreira 2021). The axiomatizations provided in (Baltag, Fiutek, and Smets 2014) used the alternative framework of Dynamic Doxastic Logic (DDL), based on "hypertheories" and "onion models" (Segerberg 1998; Segerberg 1995; Segerberg 2001). Next, (Fiutek 2013) recasted the DDL-type axioms for contraction in a DEL-framework, based on single-agent plausibility models, using standard doxastic logics such as CDL and $K\Box$ as the static base. In contrast, (Souza 2016; Souza et al. 2016; Souza, Vieira, and Moreira 2021) axiomatized single-agent contraction operations using as static base the much stronger (and arguably less intuitive, from a doxastic perspective) so-called Dynamic Preference Logic (DPL).[2]

In this paper, we build on our previous work, showing that DEL is well equipped to model several kinds of *joint* belief contraction by groups of agents. To model an act of belief contraction with ψ, we use operators of three kinds $[-^s\psi]\varphi$, $[-^c\psi]\varphi$ and $[-^m\psi]\varphi$, corresponding to three types of contraction considered in the literature: *severe withdrawal*, *conservative contraction* and *moderate contraction*. But here these are dynamic model-transformers, that change the current model by changing (some of) the agents' plausibility

[1]The distinction between static and dynamic revision makes a significant difference only for doxastic/epistemic formulas (since they may change their truth values after a revision).

[2]The label DPL in (Souza 2016) refers mainly to the account of preference logic in which operators for relation change, coming from Propositional Dynamic Logic, are used to model preference dynamics in (van Benthem and Liu 2007; Girard 2008).

relations. Moreover, they are *joint contractions* $[-_A\psi]$ (of each of the above three kinds), by which *all the agents in a given group A simultaneously perform a contraction (of the appropriate kind) with ψ* (and moreover, the fact that this is happening is common knowledge). The meaning of a construct of the form $[-^{\rho}_{A}\psi]\varphi$ (with $\rho \in \{s, c, m\}$) is that ψ will be true after group A performs a joint ρ-contraction with ψ.

We study the resulting dynamic logics, investigate their expressivity and decidability, provide some complete axiomatizations, and use them to investigate interesting properties of various forms of contraction, ending with a few concluding reflections and some open questions.

2 Background: Belief Revision in DEL

Throughout this paper, we fix a finite set of labels $\mathcal{A}$, called *agents*, and a set Φ of *atomic sentences*, representing 'ontic facts' (i.e., features of the world that do not depend on the agents' knowledge or beliefs). The semantics for the logical systems in this paper uses multi-agent epistemic plausibility models, as described in (Baltag and Smets 2008):

2.1 Knowledge and Belief in MEP-Models

A *multi-agent epistemic plausibility model* ('MEP-Model', for short) is a structure $\mathbf{S} = (S, \sim_a, \leq_a, \|\cdot\|)_{a\in\mathcal{A}}$, consisting of: a set S of *possible worlds* (or 'states'); a family of equivalence relations $\sim_a \subseteq S\times S$ (labeled by agents $a \in \mathcal{A}$), representing each agent's *epistemic indistinguishability* between possible worlds; a family of preorders (=reflexive and transitive relations) $\leq_a \subseteq S \times S$ (also labeled by agents $a \in \mathcal{A}$), representing each agent's *plausibility ranking* over worlds; and a *'valuation' function* $\|\cdot\| : \Phi \to \mathcal{P}(S)$, mapping every atomic sentence $p \in \Phi$ into a 'proposition' (sets of worlds) $\|p\| \subseteq S$, representing the worlds in which p is the case. The plausibility relations are assumed to satisfy three conditions: (1) "*plausibility implies epistemic possibility*" (i.e., $\leq_a$-comparable states are $\sim_a$-indistinguishable): $s \leq_a w$ implies $s \sim_a w$; (2) "*local connectedness*" (i.e., $\sim_a$-indistinguishable stare are $\leq_a$-comparable): $s \sim_a w$ implies either $s \leq_a w$ or $w \leq_a s$ (or both); (3) "*converse well-foundedness*": there are no infinite ascending chains $s_0 <_a s_1 <_a s_2 <_a \ldots <_a s_m < \ldots$ of 'more and more plausible' worlds; where the notation $<_a$ denotes the *strict plausibility* relation, defined by: $s <_a w$ iff $s \leq_a w$ and $w \not\leq_a s$.

Interpretation We read $s \leq_a w$ as "world w is *at least as plausible* for agent a as world s", while $s <_a w$ means that w is *'better' (more plausible)* than s for agent a. Only epistemically indistinguishable worlds (with $s \sim_a w$) can be compared by agent a: since they are both possible (given the same information), the agent must assess their relative plausibility, deciding whether w is better (i.e., $s <_a w$), or s is better (i.e., $w <_a s$), or they are equally plausible (i.e., $s \leq_a w \leq_a s$). In contrast, worlds s, w that are *epistemically distinguishable* ($s \not\sim_a w$) are *incomparable* (i.e., we have $s \not\leq_a w$ and $w \not\leq_a s$): the agent has information to tell them apart, so there is no need (or meaning) to compare their plausibility.

Propositions as sets of worlds. For any MEP-model, we can think of the subsets $P \subseteq S$ as 'propositions', that hold in a world w iff $w \in P$. The family $\mathcal{P}(S)$ of all propositions over a given MEP-model forms a Boolean algebra with the usual operations of negation/completement $\neg P := S - P$, conjunction/intersection $P \wedge Q := P \cap Q$, disjunction/union $P \vee Q := P \cup Q$, implication $P \to Q := \neg P \vee Q$ and equivalence $P \leftrightarrow Q := (P \to Q) \wedge (Q \to P)$. Semantic entailment over our model, defined as inclusion $P \subseteq Q$, structures this as a lattice, with the tautological proposition $\top := S$ as the top element, and the inconsistent proposition $\bot := \emptyset$ as the bottom element.

Maximal elements: 'most plausible' worlds. For any set $P \subseteq S$, we put

$$Max_{\leq_a} P := \{s \in P : s \not<_a w \text{ for any } w \in P\}$$

for the set of "maximal" (i.e., maximally plausible) worlds in P. The above conditions ensure that *every non-empty subset $P \subseteq S$ has maximal elements*: i.e., if $P \neq \emptyset$ then $Max_{\leq_a} P \neq \emptyset$.

Information cells and 'best' worlds. The equivalence class of a world $w \in S$ modulo the equivalence relation $\sim_a$ is denoted by

$$w(a) := \{s \in S : w \sim_a s\},$$

and called *agent a's information cell at world w*. The family $\{w(a) : s \in S\}$ of all such information cells wrt $\sim_a$ forms *agent a's information partition*. Given any non-empty subset $P \subseteq w(a)$ of an information cell $w(a)$ wrt $\sim_a$, we put

$$best_a(P) := \{s \in P : \forall t \in P(t \leq_a s\}$$

for the set of *"best worlds" in P* (according to a): i.e., those *P-worlds that are at least as plausible as all other P-worlds*. Once again, the above conditions ensure that *every non-empty subset $P \subseteq w(a)$ has 'best' worlds, and moreover these coincide with the maximal worlds in $P \cap w(a)$*: i.e., if $P \neq \emptyset$ is s.t. $P \subseteq w(a)$, then

$$best_a P = Max_{\leq_a}(P \cap w(a)) \neq \emptyset.$$

Conditional belief relations. Given a MEP-model, we can define agent-labeled *conditional belief relations* $\to^P_a \subseteq S \times S$, one for each proposition $P \subseteq S$:

$$s \to^P_a w \text{ iff } w \in best_a(P \cap s(a)) = Max_{\leq_a}(P) \cap s(a).$$

The *(unconditional) belief relations* $\to_a \subseteq S \times S$ for each agent $a \in \mathcal{A}$ can be obtained as a special case of the conditional belief relation, by taking the condition $P := \top = S$ to be the tautological proposition, so $\to_a := \to^{\top}_a$ satisfies: $s \to_a w$ iff $w \in best_a s(a) = Max_{\leq_a} s(a)$.

Epistemic and doxastic operators. Given a MEP-model, we can introduce *epistemic-doxastic operators* $K_a, \Box_a, B^P_a, B_a$ on propositions in $\mathcal{P}(S)$, defined as the Kripke modalities for respectively the indistinguishability relation $\sim_a$, the plausibility relation $\leq_a$, the conditional belief relation $\to^P_a$ and the simple belief relation $\to_a$:

$$K_a(Q) := \{s \in S : t \in Q \text{ for all } s \sim_a t\},$$

$$\Box_a(Q) := \{s \in S : t \in Q \text{ for all } s \leq_a t\},$$
$$B_a^P(Q) := \{s \in S : t \in Q \text{ for all } s \to_a^P t\},$$
$$B_a(Q) := \{s \in S : t \in Q \text{ for all } s \to_a t\},$$

for all propositions $P, Q \subseteq S$.

Interpretation: epistemic/doxastic attitudes. We read $K_a(Q)$ as saying that *agent a ('infallibly') knows* Q, $\Box_a(Q)$ as the agent *safely believes* (or "defeasibly knows") Q, $B_a^P Q$ as *the agent believes* Q *conditional on* P, and $B_a(Q)$ as *the agent (unconditionally) believes* Q. As the names point out, K_a and $\Box_a$ can be considered as forms of 'knowledge', since they are *factive*, in the sense of implying truth: $K_a(Q) \subseteq Q$ and $\Box_a(Q) \subseteq Q$. Infallible knowledge K_a is absolutely certain, fully introspective, and "irrevocable": it can never be revised. In contrast, safe belief $\Box_a$ is a 'fallible', defeasible form of knowledge, that is only positively (but not negatively) introspective (i.e., $\Box_a(Q) = \Box_a(\Box_a(Q))$ but $\neg\Box_a(Q) \not\subseteq \Box_a(\neg\Box_a(Q))$); it is 'safe' in the face of new *truthful* information, but it can be defeated (lost) if the agents receive (false) mis-information. The definition of (conditional) belief amounts to saying that Q *is believed (conditional on* P*) if it is true in all the most plausible (*P*-)worlds that are consistent with the agent's information* $s(a)$ (in the current state s). See (Baltag and Smets 2008) for more explanations and examples concerning these attitudes.

2.2 Static Logics in MEP-Models

Following (Baltag and Smets 2008), we consider two 'static' languages for multi-agent plausibility models: Conditional Doxastic Logic (CDL), and the (richer) logic $K\Box$ of (infallible) knowledge and safe belief.

Conditional Doxastic Logic. The language $\mathcal{L}_{CDL}$ of CDL is built up from atomic sentences $p \in \Phi$ using the standard Boolean constructs and the conditional belief operators $B_a^\theta\varphi$. In BNF form, the syntax is:

$$\varphi \;:=\; p \mid \neg\varphi \mid \varphi\wedge\varphi \mid B_a^\varphi\varphi$$

For any MEP-model $\mathbf{S} = (S, \sim_a, \leq_a, \|\cdot\|)_{a\in\mathcal{A}}$, the *semantics* of CDL is obtained by extending the valuation map to an *interpretation map* $\|\cdot\|_\mathbf{S} : \mathcal{L}_{CDL} \to \mathcal{P}(S)$, defined recursively: for the atomic formulas $p \in \Phi$, $\|p\|_\mathbf{S}$ is just the valuation $\|p\|$; while the recursive clauses for classical logical connectives and for the conditional-doxastic modalities use the corresponding propositional operators over $\mathcal{P}(S)$, e.g. $\|B_a^\varphi\psi\| = B_a^{\|\varphi\|}(\|\psi\|)$. The *satisfaction relation* $s \models_\mathbf{S} \varphi$ is definable as $s \in \|\varphi\|_\mathbf{S}$ (though it can also be defined independently by corresponding recursive clauses). The other classical connectives $\top, \bot, \to, \leftrightarrow$ can be defined as abbreviations in the usual way; the same goes for (unconditional) *belief* and *(infallible) knowledge*[3], which are definable in CDL as $B_a\psi := B_a^\top\psi$, $K_a\psi =: B_a^{\neg\psi}\bot$. It is also useful to introduce dual (existential) Diamond-type modalities $\langle B_a^\varphi\rangle\psi := \neg B_a^\varphi\neg\psi$, and their unconditional version $\langle B_a\rangle\psi := \langle B_a^\top\rangle\psi$, denoting (conditional and unconditional) forms of *doxastic possibility*.

[3]But defeasible knowledge (safe belief) $\Box_a$ cannot be defined in this language!

Conditional beliefs encode 'dynamic' revision Conditional belief operators $B^\varphi\psi$ provide a model for *"static" belief revision*: essentially, $B_a^\varphi\psi$ says that, after revising with φ, agent a will come to believe that ψ *was* the case (before the revision). This can be seen from the fact that, in the semantics of $B_a^\varphi\psi$, both φ and ψ are evaluated in the same (original) model: $\|B_a^\varphi\psi\|_\mathbf{S}$ is a function of $\|\varphi\|_\mathbf{S}$ and $\|\psi\|_\mathbf{S}$. In Dynamic Epistemic Logic literature, this static revision is to be distinguished from so-called *"dynamic" belief revision* (to be dealt in the next section), which concerns the revised beliefs about the state of the world *after* the revision. The distinction only makes a difference for sentences ψ that are defined using epistemic-doxastic operators, since these may change their truth values after the revision.

Axioms for CDL The following proof system CDL is sound and complete for conditional doxastic logic wrt MEP-models, cf. (Board 2002) and (Baltag and Smets 2006a):[4]

Necessitation Rule:	From $\vdash \varphi$ infer $\vdash B_a^\psi\varphi$
Normality:	$\vdash B_a^\theta(\varphi\to\psi) \to (B_a^\theta\varphi \to B_a^\theta\psi)$
Truthfulness of Knowledge:	$\vdash K_a\varphi \to \varphi$
Persistence of Knowledge:	$\vdash K_a\varphi \to B_a^\psi\varphi$
Strong Positive Introspection:	$\vdash B_a^\psi\varphi \to K_a B_a^\psi\varphi$
Strong Negative Introspection	$\vdash \neg B_a^\psi\varphi \to K_a\neg B_a^\psi\varphi$
Success:	$\vdash B_a^\varphi\varphi$
Superexpansion:	$\vdash B_a^{\varphi\wedge\psi}\theta \to B_a^\varphi(\psi\to\theta)$
Subexpansion:	$\vdash \langle B_a^\varphi\rangle\psi \to (B_a^{\varphi\wedge\psi}\theta \leftrightarrow B_a^\varphi(\psi\to\theta))$

The Logic of Knowledge and Safe Belief. In BNF form, the language $\mathcal{L}_{K\Box}$ of the logic $K\Box$ is:

$$\varphi \;:=\; p \mid \neg\varphi \mid \varphi\wedge\varphi \mid \Box_a\varphi \mid K_a\varphi$$

As before, the semantics extends the valuation in MEP-models to an interpretation map, using the operators $K_a, \Box_a$ defined in the previous section. The classical connectives are defined as abbreviations as before, and it is again useful to introduce dual (existential) Diamond-type modalities $\langle K_a\rangle\varphi := \neg K_a\neg\varphi$ and $\Diamond_a\varphi := \neg\Diamond_a\neg\varphi$, denoting forms of *epistemic possibility*. Moreover, $\mathcal{L}_{K\Box}$ is richer than $\mathcal{L}_{CDL}$: belief can now be defined as an abbreviation $B_a\varphi := \Diamond_a\Box_a\varphi$.[5] More generally, conditional belief is definable as $B_a^\varphi\psi := K_a^\varphi\Diamond_a^\varphi\Box_a^\varphi\psi$, where we are using the conditionalized versions of our modalities and their existential duals ($K_a^\varphi\psi := K_a(\varphi\to\psi)$, $\Box_a^\varphi\psi := \Box_a(\varphi\to\psi)$ and $\Diamond_a^\varphi\psi := \neg\Box_a^\varphi\neg\psi$).[6] In conclusion, *the language* $\mathcal{L}_{CDL}$ *is (translatable into) a fragment of the language* $\mathcal{L}_{K\Box}$.

Axiomatization A sound and complete axiomatization of the logic $K\Box$ was provided in (Baltag and Smets 2008). In addition to the rules and axioms of propositional logic, the *proof system* for $K\Box$ includes:[7]

[4]These axioms are in fact a modified version of the classical AGM postulates for Belief Revision, rendered in the modal language of multi-agent CDL.

[5]or equivalently, as $K_a\Diamond_a\Box_a$, which explains the above generalization to conditional belief.

[6]It is easy to see that these abbreviations match the CDL semantics of $B_a\varphi$ and $B_a^\varphi\psi$.

[7]The last axiom below is new, and is based on a simplifica-

- the Necessitation Rules for both K_a and $\Box_a$;
- the $S5$-axioms for K_a;
- the $S4$-axioms for $\Box_a$;
- $K_a\varphi \to \Box_a\varphi$;
- $(\Box_a\varphi \wedge \Diamond_a\theta) \to K_a(\Box_a\varphi \vee \Diamond_a\theta)$.

Moreover, *the logic* $K\Box$ *is decidable and has the Finite Model Property*, and thus *the logic* CDL *also has these properties* (since as we saw it can be translated into a fragment of the logic $K\Box$).

2.3 Dynamic Belief Revision and Expansion

Belief "dynamics" refers to the transformations of MEP-models that are induced by changes of the plausibility orders of some or all of the agents. The dynamical operations can correspond to forms of belief revision, expansion or contraction, as classically considered in the Belief Revision Theory literature (see e.g. (Rott 1989), (Boutilier 1996) and (Veltman 1996)). But the effect of such a model transformation goes beyond simple conditional reasoning (as in 'static' belief revision, modelled as conditional belief B^PQ), since doxastic sentences may change their truth-values

In a multi-agent setting, various forms of joint learning are discussed in (van Benthem 2007; van Benthem 2011; Baltag and Smets 2013; Baltag and Smets 2006b) in a Dynamic Epistemic Logic Framework. Here we consider two types of revision operators: *joint radical upgrade* $*^r_AP$ and *joint conservative upgrade* $*^c_AP$, for groups of agents $A \subseteq \mathcal{A}$.[8] Both are forms of (joint) group belief *revision* induced by *public announcements*: intuitively, a proposition P is *publicly announced* (to *all* agents in $\mathcal{A}$) by some (unnamed) speaker or 'source'. It is common knowledge that agents $a \in A$ belonging to some (sub)group $A \subseteq \mathcal{A}$ revise their beliefs/plausibility with P, while the outsiders $b \in \mathcal{A} - A$ keep their old plausibility order. The 'insiders' (i.e., the members of group A) trust the source of information, to some degree: they have a 'positive' attitude (being inclined to accept P), and moreover they have a *common such attitude* towards the information coming from this source. Depending on this (common) attitude, the agents $a \in A$ perform one of the two doxastic transformers (changing their plausibility relation accordingly). They perform a (joint) radical revision $*^r_AP$ when they *strongly trust* the source: in all cases (regardless of other facts), they consider that it is more plausible that the speaker tells the truth than the speaker lies. Or they do a (joint) conservative revision $*^c_AP$ when they (only) *barely trusts* the source: they believe that the speaker tells the truth, but this belief is very fragile (and can be easily overturned by further revisions). In contrast, the outsiders in $\mathcal{A} - A$ adopt a 'neutral attitude: they ignore the new information and thus "do nothing" (i.e., essentially keeping her old beliefs). Furthermore, all agents' attitudes towards the source are common knowledge: so it is common knowledge (among all agents in $\mathcal{A}$) that the insiders perform this specific doxastic upgrade, and that the outsiders keep their plausibility order unchanged.

(1) (Joint) Radical Upgrade. Given a group of agents $A \subseteq \mathcal{A}$ and a proposition $P \subseteq S$ over a MEP-model $\mathbf{S}$, the "*joint radical upgrade*" (or joint "lexicographic revision") $*^r_AP$ (also denoted in various papers by $\Uparrow_A P$) is an operation that changes any given model $\mathbf{S} = (S, \sim_a, \leq_a, \| \bullet \|)$ to a revised model $\mathbf{S} *^r_A P$ as follows: *each insider agent* $a \in A$ *"promotes" the* P*-worlds within each of her information cells, so that they become "better" (more plausible for agent* a*) than all* $\neg P$*-worlds in the same information cell*, while *keeping everything else the same* (including the relative ordering between worlds within either of the two zones (P and $\neg P$, as well as outsiders' plausibility orderings $\leq_b$ for $b \in \mathcal{A} - A$). Formally, the *radically-upgraded model* $\mathbf{S}*^r_AP = (S, \sim_a, \leq'_a, \| \bullet \|)$ has the same set of states S, the same valuation $\| \bullet \|$ and the same epistemic relations $\sim$ as the original model $\mathbf{S}$; while the new plausibility relations $\leq'$ are given, for the insiders $a \in A$, by putting: $s \leq'_a w$ iff either $s \notin P$ and $w \in P \cap s(a)$, or $s, w \in P$ and $s \leq_a w$, or else $s, w \notin P$ and $s \leq_a w$; and for the outsiders $b \notin A$, we keep the same order: $s \leq'_b w$ iff $s \leq_b w$.

(2) (Joint) Conservative Upgrade. The *"joint conservative upgrade"* (or joint "minimal revision" (Boutilier 1996)) $*^c_AP$ (denoted elsewhere by $\uparrow_A P$) performs in a sense the minimal possible revision of a model that is forced by believing the new information P. As an operation on plausibility models, it can be described as *"promoting" only the "best" (most plausible)* P*-worlds, so that they become the most plausible in their information cell*, while *keeping everything else the same*. In the joint upgrade, the insiders $a \in A$ apply this operation to their plausibility orders, while the outsiders keep their old orders. Formally, for any given MEP-model $\mathbf{S} = (S, \sim_a, \leq_a, \| \bullet \|)$, the *conservatively-upgraded model* $\mathbf{S}*^c_AP = (S, \sim_a, \leq'_a, \| \bullet \|)$ has again the same set of states S, same valuation $\|\bullet\|$ and same epistemic relation $\sim$ as the original model; while the new plausibility relations $\leq'$ are given, for the insiders $a \in A$, by putting: $s \leq'_a w$ iff either $w \in best_a(P \cap s(a))$, or else $s \leq_a w$ and $s \notin best_a(P \cap s(a))$; and for the outsiders $b \notin A$, we keep the same order: $s \leq'_b w$ iff $s \leq_b w$.

Special cases: individual and common upgrades We obtain *individual (radical or conservative) upgrades* $*^r_aP$ and $*^c_aP$ as special cases of the corresponding joint upgrades when $A = \{a\}$ is a singleton, and similarly obtain *common (radical or conservative) upgrade* $*^rP$ and $*^cP$ by taking $A = \mathcal{A}$ to be the set of all agents.

Dynamic Logics for Joint Belief Revision For both our static logics CDL and $K\Box$, we can consider their extensions $CDL*^\tau$, $K\Box*^\tau$ with dynamic operators $[*^\tau_A\varphi]\psi$, where $A \subseteq \mathcal{A}$ ranges over groups of agents and the index $\tau \in \{r, c\}$ refers to either radical or conservative upgrades.

tion of the system in (Baltag and Smets 2008), which contained instead the more complicated-looking axiom $K_a(\varphi \vee \Box_a\theta) \wedge K_a(\theta \vee \Box_a\varphi) \to K_a\varphi \vee K_a\theta$.

[8] In the DEL literature, these upgrades are usually defined for the whole group $A = \mathcal{A}$ of all agents (and denoted by $\Uparrow P$ and respectively $\uparrow P$), and considered besides other transformers, such as "joint update" (public announcement) $!P$. This last operation will play no role in our paper, as it is a form of knowledge update, rather than a belief revision operator.

The semantics is given on any MEP-model $\mathbf{S}$ by recursively extending the interpretation map $\|\bullet\|_{\mathbf{S}}$ to the new sentences, using the clause:

$$\|[*^{\tau}_{A}\varphi]\psi\|_{\mathbf{S}} = \|\varphi\|_{\mathbf{S}*^{\tau}_{A}\|\varphi\|_{\mathbf{S}}}, \quad \text{for } \tau \in \{r, c\}.$$

Joint (Radical or Conservative) Expansions As commonly understood, *belief expansion* is the 'trivial' case of belief revision, in which the new information P is consistent with the old beliefs; in this case, the new ('static') beliefs are obtained simply by putting together P and the old beliefs (and closing under logical inference): semantically, this means that the 'best' worlds after expansion with P are simply the 'old best' worlds satisfying P. But this description does not uniquely determine the effect of expansion on the 'rest' of the structure (the plausibility order between non-best worlds). And indeed, *for each of the two forms (radical and conservative) of joint dynamic belief revision*, we can also consider the corresponding forms of *(joint) expansion (radical or conservative)* $+^{r}_{A}P$ and $+^{c}_{A}P$. These are simply the *restrictions* of the (radical or conservative) upgrades to the case when *the new information P is consistent with all the beliefs of all the agents in A*. Correspondingly, we have the dynamic operators $[+^{\tau}_{A}\varphi]\psi$ for (joint) radical/conservative expansion, where $A \subseteq \mathcal{A}$ ranges over groups of agents and the index $\tau \in \{r, c\}$, with the semantics given by the following clause

$$\|[+^{\tau}_{A}\varphi]\psi\|_{\mathbf{S}} = \|pre(+^{\tau}_{A}\varphi) \to [*^{\tau}_{A}\varphi]\psi\|_{\mathbf{S}},$$

where $*^{\tau}_{A}$ is the corresponding upgrade modality and where we denoted by

$$pre(+^{\tau}_{A}\varphi) := \bigwedge_{a\in A} \langle B_a \rangle \varphi$$

the *precondition* of the joint expansions $+^{r}_{A}\varphi$ and $+^{c}_{A}\varphi$ (asserting that *no agent in group A believes that φ is false*).[9]

Before giving the specific axiomatizations for each of the logics $CDL*^{r}$, $CDL*^{c}$, $K\Box *^{r}$ and $K\Box *^{c}$, we introduce some general terminology, that allows to present these and other axiomatizations in a succinct and unified manner.

The Core Dynamic Axioms For any dynamic operator $[\dagger\varphi]\theta$ (where $\dagger$ can be radical or conservative revision, but it can also be any of the expansion or contraction operators introduced this paper), the *core dynamic axioms (and rules) for the operator* $\dagger$ are the following:

- $\dagger$-Necessitation rule: from $\vdash \theta$, infer $\vdash [\dagger\varphi]\theta$.
- $\dagger$-Distributivity (='Kripke axiom'): $[\dagger\varphi](\psi \to \theta) \to ([\dagger\varphi]\psi \to [\dagger\varphi]\theta)$.
- Atomic Preservation (=the 'Reduction axiom for atoms'): $[\dagger\varphi]p \leftrightarrow p$.
- Determinism (=the 'Reduction axiom for negation'): $[\dagger\varphi]\neg\theta \leftrightarrow \neg[\dagger\varphi]\theta$.

[9]Note that (radical and conservative) expansion operators are obviously definable as abbreviations in the dynamic languages for (radical and conservative) upgrades/revision.

Note that, using these core axioms, one can *prove* the so-called 'Reduction axiom for conjunction": $[\dagger\varphi](\psi \wedge \theta) \to ([\dagger\varphi]\psi \wedge [\dagger\varphi]\theta)$.

Axiomatizations of $CDL*^{\tau}$, with $\tau \in \{r, c\}$. We can obtain sound and complete axiomatizations for these dynamic logics, by taking together the following: (1) the axioms and rules of the static logic CDL; (2) the above core dynamic axioms for $*^{\tau}_{A}$; (3) the appropriate version (depending on whether $\tau = r$ or $\tau = c$) of the following "Reduction axioms for conditional belief" for all *insiders* $a \in A \subseteq \mathcal{A}$:

$$\begin{aligned}
[*^{r}_{A}\varphi]B^{\psi}_{a}\theta \quad &\longleftrightarrow \quad B^{\varphi\wedge[*^{r}_{A}\varphi]\psi}_{a}[*^{r}_{A}\varphi]\theta \wedge \\
&\qquad (K^{\varphi}_{a}[*^{r}_{A}\varphi]\neg\psi \to B^{[*^{r}_{A}\varphi]\psi}_{a}[*^{r}_{A}\varphi]\theta),\\
[*^{c}_{A}\varphi]B^{\psi}_{a}\theta \quad &\longleftrightarrow \quad B^{\varphi}_{a}([*^{c}_{A}\varphi]\psi \to [*^{c}_{A}\varphi]\theta) \wedge \\
&\qquad (B^{\varphi}_{a}[*^{c}_{A}\varphi]\neg\psi \to B^{[*^{c}_{A}\varphi]\psi}_{a}[*^{c}_{A}\varphi]\theta),
\end{aligned}$$

where we used the abbreviation $K^{\varphi}_{a}\psi := K_a(\varphi \to \psi)$; (4) for the 'outsiders' $b \notin A$, the Reduction axioms for conditional belief are simpler (for both $\tau \in \{c, r\}$):

$$[*^{\tau}_{A}\varphi]B^{\psi}_{b}\theta \leftrightarrow B^{[*^{\tau}_{A}\varphi]\psi}_{b}[*^{\tau}_{A}\varphi]\theta.$$

Moreover, these axioms can be used to show that *the dynamic logics $CDL*^{\tau}$ are co-expressive with their static base CDL* (and thus are decidable).

Remarks There exists an analogue complete axiomatization for the logic $K\Box *^{r}$, using reduction axioms. But for $K\Box *^{c}$, there is no reduction law for $\Box_a$ after a conservative upgrade $*^{c}_{A}\varphi$: in order to obtain a complete axiomatization, one needs to move to a richer static base language. We skip the details in this short proceedings paper.

3 Dynamic Belief Contraction for Groups and Individuals

We now proceed to study various *dynamic belief contraction* operations, corresponding to various forms of "suspending judgment". Each gives rise to extensions of the above languages with corresponding dynamic operators. We completely axiomatize the resulting dynamic logics using reduction axioms, and we study their relations with the dynamic revision and expansion operations considered above.

3.1 Three Forms of Contraction and Their Logics

In this paper, we consider three operations of *(joint) dynamic belief contraction* with a proposition P (by a group $A \subseteq \mathcal{A}$): joint *severe withdrawal* $-^{s}_{A}P$, joint *conservative contraction* $-^{c}_{A}P$, and joint *moderate contraction* $-^{m}_{A}P$. Each of these originates in the Belief Revision Theory, where they have been considered only in a 'static', single-agent context, while here we consider them as dynamic transformations of multi-agent plausibility models.

Joint Severe Withdrawal In its dynamic, multi-agent form, the (joint) "severe withdrawal" $-^{s}_{A}P$ (cf. (Pagnucco and Rott 1999) for this terminology, though also known as "mild contraction" (Levi 2004) or "Rott contraction" in (Ferme and Rodriguez 1998)), is an operation on MEP-models that changes any model as follows: for each 'insider' $a \in A$, *the worlds that are at least as plausible ($\geq_a$) as the 'best' $\neg P$-worlds are equally "promoted"* above all others, so they all *become 'the best'* (most plausible) worlds within each information cell), while in rest everything else stays the same. Formally, for a given MEP-model $\mathbf{S} = (S, \sim_a, \leq_a, \|\bullet\|)$ and a given proposition $P \subseteq S$, the result of joint severe withdrawal with P by a group A is a model $\mathbf{S}-^s_A P = (S, \sim_a, \leq'_a, \| \bullet \|)$, having again the same set of states S, the same valuation $\| \bullet \|$ and the same epistemic relations $\sim$ as the original model $\mathbf{S}$; while the new plausibility relations $\leq'$ are given, for the insiders $a \in A$, by putting: $s \leq'_a w$ iff either $s \leq_a w$, or else $t \leq_a w$ for some $t \in best_a(s(a) \setminus P)$; and for the outsiders $b \notin A$, we keep the same order: $s \leq'_b w$ iff $s \leq_b w$.

Joint Conservative Contraction The dynamic operation of (joint) "conservative contraction" $-^c_A P$ (also called "partial meet contraction" by Arlo-Costa, and "AGM contraction" by Rott) changes any MEP-model as follows: for each 'insider' $a \in A$, *the 'best' $\neg P$-worlds are 'promoted' to become equally plausible to the currently 'best' worlds* (in the same information cell), while in rest everything else stays the same. Formally, for a given MEP-model $\mathbf{S} = (S, \sim_a, \leq_a, \|\bullet\|)$ and a given proposition $P \subseteq S$, the result of joint conservative contraction with P by a group A is a model $\mathbf{S}-^c_A P = (S, \sim_a, \leq'_a, \| \bullet \|)$, having again the same set of states S, the same valuation $\| \bullet \|$ and the same epistemic relations $\sim$ as the original model $\mathbf{S}$; while the new plausibility relations $\leq'$ are given, for the insiders $a \in A$, by putting: $s \leq'_a w$ iff either $w \in best_a(s(a) \setminus P)$, or else $s \leq_a w$ and $s \notin best_a(s(a) \setminus P)$; and for the outsiders $b \notin A$, we keep the same order: $s \leq'_b w$ iff $s \leq_b w$.

Joint Moderate Contraction This is a dynamic operation $-^m_A P$ that changes any MEP-model as follows: for each 'insider' $a \in A$, *the 'best' $\neg P$-worlds are 'promoted' to become equi-plausible to the currently 'best' worlds* (in the same information cell), then *next all the remaining $\neg P$-worlds become better than the remaining P-worlds* (in the same information cell), while *within these two zones the old ordering remains the same*; and in rest everything else stays the same. Formally, for a given MEP-model $\mathbf{S} = (S, \sim_a, \leq_a, \| \bullet \|)$ and a given proposition $P \subseteq S$, the result of joint moderate contraction with P by a group A is a model $\mathbf{S}-^m_A P = (S, \sim_a, \leq'_a, \| \bullet \|)$, having again the same set of states S, the same valuation $\| \bullet \|$ and the same epistemic relations $\sim$ as the original model $\mathbf{S}$; while the new plausibility relations $\leq'$ are given, for the insiders $a \in A$, by putting $s \leq'_a w$ iff: (1) either $w \in best_a(s(a) \setminus P)$; or (2) $s \in P$ and $w \in s(a) \setminus P$; or (3) $s, w \in P$ and $s \leq_a w$; or (4) $s, w \notin P$ and $s \leq_a w$; while for the outsiders $b \notin A$, we keep the same order: $s \leq'_b w$ iff $s \leq_b w$.

Dynamic logics of contraction, revision and expansion We can extend all the above languages in a natural way with dynamic belief operators $[-^\rho_A \varphi]\psi$, where $A \subseteq \mathcal{A}$ ranges over groups of agents and the index $\rho \in \{s, c, m\}$ refers to either severe withdrawal, or conservative contraction, or moderate contraction. Indeed, we can add any of these to the static languages CDL or $K\Box$, obtaining dynamic logics of the form $CDL-^\rho$ or $K\Box-^\rho$ (with $\rho \in \{s, c, m\}$); or we can even mix and match them with either of the dynamic revision operators above, obtaining logics of the form $CDL *^\tau -^\rho$ or $K\Box *^\tau -^\rho$ (with $\tau \in \{r, c\}$ and $\rho \in \{s, c, m\}$). The semantics of the dynamic contraction operators is given on any MEP-model $\mathbf{S}$ by recursively extending the interpretation map $\| \bullet \|_{\mathbf{S}}$ to the new sentences, using the clause:

$$\|[-^\rho_A \varphi]\psi\|_{\mathbf{S}} = \|\varphi\|_{\mathbf{S}-^\rho_A \|\varphi\|_{\mathbf{S}}}, \quad \text{for } \rho \in \{s, c, m\}.$$

Equivalent representation: Grove spheres Multi-agent plausibility models can be equivalently represented as *'Grove sphere' models* (Grove 1988): for every agent a, each of her information cells $s(a)$ at any state s is the union $\bigcup \mathcal{S}_a(s)$ of a family $\mathcal{S}_a(s)$ of sets, called 'spheres'. The spheres in each $\mathcal{S}_a(s)$ are assumed to be non-empty and well-ordered wrt inclusion. The agent believes Q given P at state s if the smallest sphere in $\mathcal{S}_a(s)$ that has non-empty intersection with P is included in Q. In a plausibility model, the spheres correspond to the sets that are upwards-closed wrt $\leq_a$.[10] But Grove spheres have their own independent motivation, as representing the *primary evidence* or "core beliefs" that support an agent's beliefs: Q is believed (conditional on P) if the agent has evidence (consistent with P) that entails Q.

3.2 Features and Bugs

We now look at various (positive and negative) features of the contraction operators, and the way they can be expressed in our modal languages.

Dynamic versions of Levi identity The so-called Levi identity is a principle in classical Belief Revision Theory, stating that a 'good' belief revision operation with P can be obtained by first applying a contraction with $\neg P$, then performing an expansion with P. In our framework, we say that *a given contraction operation $\rho \in \{s, c, m\}$ satisfies dynamic Levi identity with respect to a given expansion/revision operation $\tau \in \{r, c\}$* if the following schema is valid (for all groups $A \subseteq \mathcal{A}$ and all atomic sentences p):

$$[*^\tau_A p]\theta \leftrightarrow [-^\rho_A \neg p][+^\tau_A p]\theta.$$

Let us call this schema *'dynamic' Levi*$^{\rho,\tau}$, for short. A weaker, "more static" version of Levy identity can be obtained by replacing dynamic revision by static revision (conditional belief), and replace θ by any specific belief Bq. We say that *a given contraction operation $\rho \in \{s, c, m\}$, satisfies static Levi identity with respect to a given expansion/revision operation $\tau \in \{r, c\}$* if the following schema is valid (for all $a \in A \subseteq \mathcal{A}$ and all atomic sentences p, q):

$$[*^\tau_A p]B_a q \leftrightarrow [-^\rho_A \neg p]B^p_a q.$$

We call this later schema *'static' Levi*$^{\rho,\tau}$. Using the equivalence between $[*^\tau_A p]B_a q$ and $B^p_a q$ (for $\tau \in \{r, c\}$), it is easy

[10] And vice-versa, given a sphere model, we can recover the plausibility relations by putting: $s \leq_a w$ if $s \sim_a w$ and $\forall S' \in \mathcal{S}_a(s)\, (s \in S' \Rightarrow w \in S')$.

to see that *the dynamic Levy identity entails the static Levi identity* (but not the other way around).

Observations *Severe withdrawal fails the static (hence, also the dynamic) Levi identity wrt both forms of revision*; i.e., (both 'static' and 'dynamic') $Levi^{s,r}$ and $Levi^{s,c}$ schemas are invalid. In contrast, *both conservative contraction and minimal contraction validate the 'dynamic' (hence, also the static) Levi identity wrt both forms of revision*; i.e., (both 'static' and 'dynamic' forms of) the schemas $Levi^{c,r}$, $Levi^{c,c}$, $Levi^{m,r}$ and $Levi^{m,c}$ are valid.

Proviso One may be tempted to consider an even "more dynamic" form of Levi identity, obtained by replacing the atomic sentence p by any arbitrary formula φ, i.e. requiring the validity of the schema $[*^{\tau}_{A}\varphi]\theta \leftrightarrow [-^{\rho}_{A}\neg\varphi][+^{\tau}_{A}\varphi]\theta$, or at least of the weaker schema $[*^{\tau}_{A}\varphi]B_a q \leftrightarrow [-^{\rho}_{A}\neg\varphi]B^{\varphi}_{a}q$ (for all atoms q). But the fact that doxastic formulas φ may change their truth values after the contraction suggests that these forms do not really match their intended meaning; and indeed, it is easy to see that *neither of these schemas is valid* for any pair (ρ, τ) (-except for non-doxastic formulas φ, that are simple Boolean combinations of atoms).

Expulsiveness As we saw, severe withdrawal fails all forms of Levi identity. But many authors consider severe withdrawal to be a 'bad' candidate for modeling contraction, because of a deeper reason. Namely, it *does* satisfy a highly implausible property, called *Expulsiveness*: essentially, this states that any two (even unrelated) beliefs will always be disturbed by each other's contraction! In our framework, Expusiveness is embodied by the following valid schema (for all $a \in A \subseteq \mathcal{A}$ and atomic sentences p, q):

$$(\neg B_a p \wedge \neg_a Bq) \rightarrow ([-^{s}_{A}p]B_a q \vee [-^{s}_{A}q]B_a p).$$

Recovery Postulate and its failures. Severe withdrawal has also been criticised for failing to satisfy the so-called Recovery Postulate (see (Arlo Costa and Pedersen 2011)). In the syntactic setting of classical AGM Theory, the Recovery Postulate requires that after we contract with P and expand again with P, we reinstate the original theory. In our framework, we can again distinguish between various versions of this principle. We say that *a contraction operation $\rho \in \{s, c, m\}$ satisfies the dynamic Recovery Postulate wrt an expansion/revision operation $\tau \in \{r, c\}$* if the following schema is valid (for all groups $A \subseteq \mathcal{A}$ and all atoms p):

$$\theta \leftrightarrow [-^{\rho}_{A}p][+^{\tau}_{A}p]\theta.$$

The contraction operation satisfies the *static Recovery Postulate* wrt an expansion/revision operation $\tau \in \{r, c\}$ if the following schema is valid (for all $a \in A$ and all atoms p, q):

$$B_a q \leftrightarrow [-^{\rho}_{A}p]B^{p}_{a}q$$

Using again the equivalence between $[*^{\tau}_{A}p]B_a q$ and $B^{p}_{a}q$ (for $\tau \in \{r, c\}$), it is easy to see that *dynamic Recovery entails the static Recovery* (but not the other way around).

Observations *Severe withdrawal fails the static (hence, also the dynamic) Recovery Postulate wrt both forms of revision* (radical and conservative). See Section 3.4 for a counterexample. In contrast, *both conservative contraction and minimal contraction validate the dynamic (hence, also the static) Levi identity wrt both forms of revision.* [11]

3.3 Why These Three Choices and No Others?

Readers who are familiar with the puzzling plethora of contraction operations in the Belief Revision Literature may question our focus on these particular three (severe withdrawal, conservative contraction and moderate contraction), and may ask why are these more interesting than other choices, such as maxi-choice or full meet contraction. To this, we answer that first, this paper is more of a sample conceptual exercise. Since we cannot possibly aim to axiomatize all the many contraction functions proposed in the literature in a 10-page paper, we chose three that allow for DEL-style reductions with respect to well-known and natural static doxastic logics such as CDL or $K\Box$.

But second, we think that our three choices are of particular interest. Concerning several withdrawal, a number of authors such as Rott and Levi consider this to be the most natural form of contraction. Pagnucco and Rott (1999) provide a strong argument in favor of 'severe withdrawal' in terms of rationality postulates: the AGM Principle of Informational Economy "cannot be given unrestrained prominence over other rationality principles". When contracting a belief, it makes sense to treat this principle on the same footing as the principles of indifference and preference, which Pagnucco and Rott describe as 'If one object is held in equal or higher regard than another, the former should be treated no worse than the latter'. Pagnucco and Rott argue that these principles do some work in AGM partial-meet contraction, but point to 'severe withdrawal', as an operation that more faithfully adheres to them.

The choice of terminology may lead to the impression that this operation "is too drastic" and "removes everything" of interest from the doxastic structure. But this is simply not the case: indeed, Levi (Levi 2004) criticised the use of the term 'severe withdrawal' and proposed instead the term 'mild contraction'. To see how natural is this operation from an evidential perspective, it is best to look at the equivalent formulation of plausibility models in terms of Grove spheres. If we think of the spheres as the "evidence" or "core beliefs" that support an agent's beliefs, then severe withdrawal with P simply *removes all the evidence for P*, while keeping all other evidence. This seems the most natural way of withdrawing or 'ignoring' all evidence for P (say, because it is declared 'inadmissible').

It may seem that the failure of the Levi identity and of Recovery Principle for severe withdrawal constitutes a strong objection against this contraction operator. But equally, the failure of these principles can be turned into an argument against *them*. Indeed, many authors argued against these postulates, and they seem particular inadequate when dealing with evidence withdrawal. Suppose that in a trial one

[11] As before, there are other, "more dynamic" forms of Recovery, embodied by the schemas $\theta \leftrightarrow [-^{\rho}_{A}\varphi][+^{\tau}_{A}\varphi]\theta$ and $B_a q \leftrightarrow [-^{\rho}_{A}\varphi]B^{p}_{a}q$, which fail for all dynamic contraction operations.

is first instructed to dismiss all existing evidence for P (as inadmissible, unwarranted, illegally obtained, etc), but then some new testimony in favor of P is added (as fresh evidence). Clearly, not all the original beliefs are automatically reinstated after this! The old evidence (which may have been abundant, corresponding to several Grove spheres, supporting many other beliefs besides P) still remains inadmissible, but instead we have one new piece of evidence for P. Thus, many old beliefs may still be lost, although the belief in P itself is restored: the Recovery Principle is obviously untenable in this interpretation!

In our view, the strongest objection against severe withdrawal is the phenomenon of Expulsiveness. But this feature is caused by the "local-connectedness" condition on plausibility models (-every two worlds in the same information cell are comparable in plausibility). This is a 'hyper-rationality' postulate, meant to ensure that belief revision satisfies all the AGM postulates (or that conditional belief satisfies Rational Monotonicity). Many authors expressed doubts about this assumption, and dropped it (-in which case conditional beliefs only satisfy Cautious Monotonicity).[12] It would be interesting to axiomatize severe withdrawal in generalized settings, eg., van Benthem's and Pacuit's evidence models (van Benthem and Pacuit 2011), where the natural plausibility induced by (non-nested) evidence will not necessarily satisfy local-connectedness.

Moving on to our other two choices (conservative and moderate contraction), they both belong to the large class of "partial-meet contractions", which includes other salient interesting operations considered in the literature, e.g., *maxi-choice* and *full-meet contraction*. But these later ones are "extreme", limiting cases. As (Hansson 2022) mentions: "Both limiting cases (maxi-choice and full-meet) are technically useful, but they also both have highly unrealistic properties." Maxi-choice contraction by P proceeds non-deterministically, by randomly picking *one* world from the most plausible non-P worlds and adding it to the smallest sphere: the choice seems rather arbitrary and unjustified. This non-deterministic nature will also pose problems to DEL-style axiomatization: one cannot have Reduction axioms, since the choice is not pre-determined by the doxastic structure (and the proposition P). On the other hand, full-meet contraction adds *all* the non-P worlds to the smallest sphere (totally disregarding the plausibility ranking between non-P worlds). This is often considered as too radical in the opposite sense: it removes too many of the original beliefs.

In contrast, conservative contraction seems a much more natural choice, designed to retain as much of the original beliefs as possible, while giving up the belief in P (in a deterministic, non-arbitrary way): only (and all) the most plausible non-P worlds are added to the smallest sphere. According to some authors: "Conservative contraction is designed for practical belief base modification, whereas full meet is often used as a theoretical "building block" or a point of reference because of its purely logical nature.". Finally, moderate contraction is an interesting and practical combination of conservative contraction and lexicographic revision with $\neg P$: it does the same as conservative revision as far as the new smallest sphere is concerned; but in rest, it gives relative priority to the non-P worlds over the P-worlds (which are "demoted", while still keeping the relative ranking within each zone, outside the smallest sphere). This is more "radical" than conservative-contraction (giving up more of the original beliefs), but it maintains more of the original structure than full-meet contraction.

3.4 Examples

Let us now consider some examples, to illustrate the differences between the various joint contraction operations.

Example 1. Let $A = \{a, m\}$ be the set of two agents consisting of Professors Albert Winestein (a) and Mary Curry (m). Albert feels that he is a genius (denoted by g). He knows that there are only two possible explanations for this feeling: either he is a genius g or he is drunk d, i.e. $K_a(g \vee d)$. Albert doesn't feel drunk, so he believes that he is a sober genius, i.e. $B_a(\neg d \wedge g)$. However, if Albert realized that he's drunk, he'd think that his genius feeling was just the effect of the drink; i.e. after learning he is drunk he would come to believe that he was just a drunk non-genius $B_a^d \neg g$. In reality though, he is both drunk and a genius (i.e. $d \wedge g$). Albert's friend, Mary, obtains some evidence that indicates that Albert is a drunk (she sees the result of a blood test) and evidence for being a genius (she hears a testimony about his scientific work). As a result, she believes him to be a drunk genius $B_m(d \wedge g)$. The weight of her evidence for him being drunk or a genius is equally strong so that if she were to drop either one of the pieces of evidence, she would consider $d \wedge \neg g$ and $\neg d \wedge g$ to be equally plausible. Albert knows that Mary believes that he is a drunk genius, i.e., $K_a B_m(d \wedge g)$. Moreover, we take all the agents' epistemic and doxastic attitudes to be common knowledge in the group of Albert and Mary.

Fig.1. provides the MEP model for this scenario, where the arrows depict the agent-labeled plausibility relations.[13] All the statements of example 1 are true at the real world s (colored in grey) in the given model.

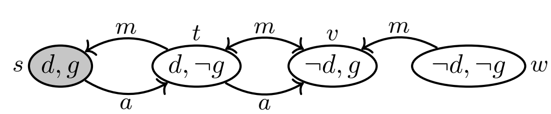

Figure 1: MEP model M for Example 1.

Example 1 Continued. Next, both agents learn that the blood test was contaminated and is not reliable. As a consequence, Mary suspends her judgement on whether or not Albert is drunk. Note that Albert already believed $\neg d$ ($M, s \models$

[12]In terms of Grove spheres, local connectedness is the Nestedness assumption (-the spheres are 'nested', i.e., totally ordered by inclusion). But this condition does not seem natural when adopting an "evidential" interpretation of the spheres: there is no reason to assume that the evidence comes in this neat decreasing-strength order.

[13]In all the drawings in this paper, we omit the reflexive loops, and often also omit arrows that can be obtained by transitivity.

$B_a\neg d$ in Fig.1), so the joint contraction with d does not affect Albert's first-order beliefs (but only his higher-order beliefs about Mary's opinion). What happens to Mary's *conditional* beliefs depends on the type of contraction operation. Fig.2. represents the result of joint severe withdrawal with d; Fig.3 represents the situation after joint conservative contraction with d; while Fig.4 is the scenario after joint moderate contraction with d.

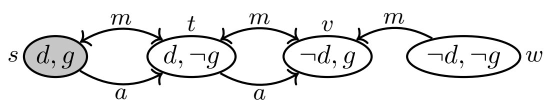

Figure 2: MEP model after the joint act of severe withdrawal with d performed on M.

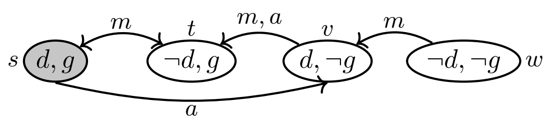

Figure 3: MEP model after the joint act of conservative contraction with d performed on M.

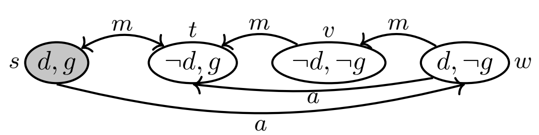

Figure 4: MEP model after the joint act of moderate contraction with d performed on M.

Discussion. After Mary performs a contraction by d, she suspends her judgement in d. Indeed, the models in Fig.2,3,4 satisfy $\neg B_m d \wedge \neg B_m \neg d$. However the *conditional* beliefs of Mary differ, depending on the specific type of contraction: both after severe withdrawal or conservative contraction with d in model M we'll get $B_m^{\neg g} d$ (i.e., we have $M \models [-^s_A d]B_m^{\neg g} d$ and $M \models [-^c_A d]B_m^{\neg g} d$); but *not* after the moderate contraction with d (since $M \models [-^m_A d]\neg B_m^{\neg g} d$). Moreover, conditional on d being the case, Mary would regain her previous belief in Albert being a genius in the models of Fig.3 and Fig.4. This is not the case in Fig.2, which illustrates the failure of Recovery Postulate for severe withdrawal.

A counterexample to Levi Identity. Fig.5 presents the model in which Mary has performed a (radical) belief expansion with $\neg d$ after the joint severe withdrawal of d (here the input model is Fig.2). Observe that Fig.6 presents the model in which both agents radically revise their beliefs with $\neg d$ directly in model M (here the input model is Fig.1). We observe that the act of a severe withdrawal followed by a radical expansion versus a direct radical revision leads to different results. In particular, Mary's conditional beliefs in the models of Fig.5 and Fig.6 are different as illustrated by the failures of the dynamic Levi identity $M, s \models [-^s_A d][+^r_A \neg d]\neg B^d_m g$ while $M, s \models [*^r_A \neg d]B^d_m g$.

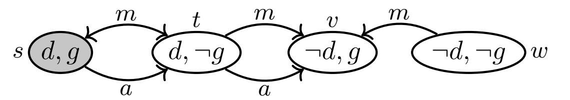

Figure 5: MEP model after the joint act of severe withdrawal of d on model M (i.e. $M -^s_A d$) followed by the joint radical expansion with $\neg d$ (i.e. $+^r_A \neg d$).

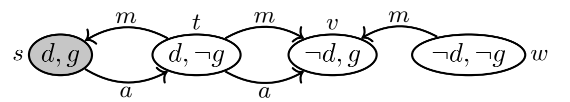

Figure 6: MEP model after the joint radical revision with $\neg d$ performed on M (i.e. $M *^r_A \neg d$).

4 Complete Axiomatizations

In this section, we give sound and complete axiomatizations of some of the above dynamic logics, via DEL-style Reduction Axioms. We also show that these logics are co-expressive with their static counterparts, and are thus decidable and have the finite model property.

Proposition 1. *A sound and complete proof system for the logic $K\Box-^s$ of severe withdrawal (over the static language $\mathcal{L}_{K\Box}$, in which the belief contraction operator of a formula φ is taken to be $[-^s_A\varphi]$) is obtained by putting together the following: (1) the axioms and rules of the complete proof system for $K\Box$; (2) the core dynamic axioms and rules for $-^s_A$; the following "Reduction axioms for Knowledge and Safe Belief" for* insiders $a \in A$:

$$[-^s_A\varphi]K_a\theta \leftrightarrow K_a[-^s_A\varphi]\theta,$$

$$[-^s_A\varphi]\Box_a\theta \leftrightarrow (\Box_a[-^s_A\varphi]\theta \wedge (\neg K_a\varphi \to \langle K_a\rangle(\neg\varphi \wedge \Box_a[-^s_A\varphi]\theta)));$$

while for outsiders $b \notin A$*, the Reduction Laws are simpler:*

$$[-^s_A\varphi]K_b\theta \leftrightarrow K_b[-^s_A\varphi]\theta, \qquad [-^s_A\varphi]\Box_b\theta \leftrightarrow \Box_b[-^s_A\varphi]\theta.$$

Proposition 2. *A sound and complete proof system for the logic $CDL-^c$ of conservative contraction over the static language $\mathcal{L}_{CDL}$ is obtained by putting together the following: (1) the axioms and rules of the complete proof system for CDL; (2) the core dynamic axioms and rules for $-^c_A$; (3) the following reduction axioms for* insiders $a \in A$:

$$\begin{aligned}[-^c_A\varphi]B^\psi_a\theta \quad \Longleftrightarrow \quad & B_a([-^c_A\varphi]\psi \to [-^c_A\varphi]\theta) \\ & \wedge B^{\neg\varphi}_a([-^c_A\varphi]\psi \to [-^c_A\varphi]\theta) \\ & \wedge \left(B^{\neg\varphi}_a[-^c_A\varphi]\neg\psi \to B^{[-^c_A\varphi]\psi}_a[-^c_A\varphi]\theta\right);\end{aligned}$$

while for outsiders $b \notin A$*, the Reduction Laws are simpler:*

$$[-^c_A\varphi]B^\psi_b\theta \leftrightarrow B^{[-^c_A\varphi]\psi}_b[-^c_A\varphi]\theta.$$

Proposition 3. *A sound and complete proof system for the logic $CDL{-}^m$ of moderate contraction over the static language $\mathcal{L}_{CDL}$ is obtained by putting together the following: (1) the axioms and rules of the complete proof system for CDL; (2) the core dynamic axioms and rules for $-^m_A$; (3) the following reduction axioms for* insiders $a \in A$:

$$\begin{aligned}[-^m_A\varphi]B^\psi_a\theta \quad\longleftrightarrow\quad & B_a([-^m_A\varphi]\psi \to [-^m_A\varphi]\theta)\\ & \wedge B_a^{\neg\varphi\wedge[-^m_A\varphi]\psi}[-^m_A\varphi]\theta\\ & \wedge\left(K_a^{\neg\varphi}[-^m_A\varphi]\neg\psi \to B_a^{[-^m_A\varphi]\psi}[-^m_A\varphi]\theta\right);\end{aligned}$$

while for outsiders $b \notin A$, *the Reduction Laws are simpler:*

$$[-^m_A\varphi]B^\psi_b\theta \longleftrightarrow B_b^{[-^m_A\varphi]\psi}[-^m_A\varphi]\theta.$$

An immediate consequence is the following:

Proposition 4. *The dynamic logics $K\Box{-}^s$, $CDL{-}^c$, and respectively $CDL{-}^m$ are co-expressive with their static base logics, and thus are decidable.*

5 Conclusions and Future Work

In this paper we introduced three types of joint contraction operations, investigated their properties, and correspondingly provided complete axiomatisations for their dynamic logics. We illustrated the differences between the different contraction operators with a number of examples, and studied the expressivity of the resulting dynamic logics.

Note that we did *not* provide axiomatizations for the logic of severe contraction over the static language $\mathcal{L}_{CDL}$, nor for the logic moderate contraction over the static language base $\mathcal{L}_{K\Box}$. The problems of axiomatizing the logics $CDL{-}^s$, $K\Box{-}^c$ and $K\Box{-}^m$ (or some natural extensions of these) are still open. Both of them seem challenging questions. The most natural options seem to extend the static base with "contracted conditional beliefs" $B_a^{-\phi}\psi$ or contracted safe beliefs $\Box_a^{-\phi}\psi$, that encode forms of "static" belief contraction (similarly to the way conditional beliefs $B^\phi\psi$ encode static belief revision). These would allow us to write reduction laws for the corresponding dynamic operators.

Another open problem is to axiomatize (and provide reduction axioms for) contraction in the presence of common knowledge, common belief, and other group attitudes. Since joint contractions affect the group's common beliefs, this would be a very natural extension. However, this could be a tough challenge. Dynamic reduction laws for common knowledge are notoriously difficult to have, in some cases impossible; and when they do exist, they can be too complicated to be of any practical use. Still, such work would be interesting, at least as a proof of concept.

Finally, the hardest open problem of all is to develop a unified account for multi-agent contraction and revision, e.g. one that extends the general Action-Priority Update mechanism for interactive belief revision (Baltag and Smets 2008).

We leave the investigation of these and other questions for future work.

AI Declaration

The authors have not employed any Generative AI tools.

References

Alchourron, C.; Gardenfors, P.; and Makinson, D. 1985. On the logic of theory change: Partial meet contraction and revision functions. *The Journal of Symbolic Logic* 50(2):510–530.

Arlo Costa, H., and Pedersen, A. P. 2011. Belief revision. In *The Continuum Companion to Philosophical Logic*. Continuum International Publishing Group. 450–502.

Baltag, A., and Moss, L. S. 2004. Logics for epistemic programs. *Synthese* 139(2):165–224.

Baltag, A., and Smets, S. 2006a. Conditional doxastic models: a qualitative approach to dynamic belief revision. In *Electronic Notes in Theoretical Computer Science*, volume 165. 5–21.

Baltag, A., and Smets, S. 2006b. Dynamic belief revision over multi-agent plausibility models. In G. Bonanno, W. van der Hoek, M. W., ed., *Proceedings of LOFT 2006, 11-24. University of Liverpool*.

Baltag, A., and Smets, S. 2006c. Dynamic belief revision over multi-agent plausibility models. In G. Bonanno, W. van der Hoek, M. W., ed., *Pre-proceedings of the 7th Conference on Logic and the Foundations of Game and Decision Theory (LOFT 2006)*.

Baltag, A., and Smets, S. 2008. A qualitative theory of dynamic interactive belief revision. In Bonanno, G.; van der Hoek, W.; and Wooldridge, M., eds., *Texts in Logic and Games*, volume 3. Amsterdam University Press. 13–60.

Baltag, A., and Smets, S. 2013. Protocols for belief merge: Reaching agreement via communication. *Logic Journal of the IGPL* 21(3):468–487.

Baltag, A.; Fiutek, V.; and Smets, S. 2014. DDL as an 'internalization' of dynamic belief revision. In *Krister Segerberg on Logic of Actions. Outstanding Contributions to Logic*, volume 1. Springer.

Baltag, A.; Moss, L.; and Solecki, S. 1998. The logic of public announcements, common knowledge and private suspicions. In *Proceedings of TARK'98 (Seventh Conference on Theoretical Aspects of Rationality and Knowledge)*. Morgan Kaufmann Publishers. 43–56.

Board, O. 2002. Dynamic interactive epistemology. *Games and Economic Behaviour* 49:49–80.

Boutilier, C. 1996. Iterated revision and minimal change of conditional beliefs. *Journal of Philosophical Logic* 25(3):262–305.

Ferme, E., and Rodriguez, R. 1998. A brief note about Rott contraction. *Logic Journal of the IGPL* 6:835–842.

Fiutek, V. 2013. *Playing with knowledge and belief*. ILLC Dissertation Series DS-2013-02, University of Amsterdam.

Girard, P. 2008. Modal logic for belief and preference change. In *Ph.D. thesis, Stanford University*.

Grove, A. 1988. Two modellings for theory change. *Journal of Philosophical Logic* 17:157–170.

Hansson, S. 2022. Logic of belief revision. *The Stanford Encyclopedia of Philosophy (Spring 2022 Edition)*.

Levi, I. 2004. *Mild Contraction: Evaluating Loss of Information due to Loss of Belief*. Oxford University Press Oxford.

Pagnucco, M., and Rott, H. 1999. Severe withdrawal - and recovery. *Journal of Philosophical Logic* 28 (corrected reprint in Feb 2000):501–547.

Rott, H. 1989. Conditionals and theory change: revisions, expansions, and additions. *Synthese* 81(1):91—113.

Segerberg, K. 1995. Belief revision from the point of view of doxastic logic. *Bulletin of the IGPL* 3.

Segerberg, K. 1998. Irrevocable belief revision in dynamic doxastic logic. *Notre Dame Journal of Formal Logic* 39(3).

Segerberg, K. 2001. A completeness proof in full ddl. *Logic and Logical Philosophy* 9.

Souza, M.; Moreira, A.; Vieira, R.; and Meyer, J. 2016. Erratum to: Priorities and preferences: A study based on contraction. In *Proceedings of 15th Int. Conf. on Principles of Knowledge Representation and Reasoning (KR 2016)*.

Souza, M.; Vieira, R.; and Moreira, A. 2021. Dynamic preference logic meets iterated belief change: Representation results and postulates characterization. *Theoretical Computer Science* 872:15–40.

Souza, M. 2016. Choices that make you change your mind: a dynamic epistemic logic approach to the semantics of bdi agent programming languages. In *Ph.D. dissertation, Universidade Federal do Rio Grande do Sul*.

van Benthem, J., and Liu, F. 2007. Dynamic logic of preference upgrade. *Journal of Applied Non-Classical Logics* 17:2:157–182.

van Benthem, J., and Pacuit, E. 2011. Dynamic logics of evidence-based beliefs. *Studia Logica* 99(61):61–92.

van Benthem, J. 2007. Dynamic logic of belief revision. *Journal of Applied Non-Classical Logics* 17(2):129—-155.

van Benthem, J. 2011. *Logical Dynamics of Information and Interaction*. Cambridge University Press.

van Ditmarsch, H.; van der Hoek, W.; and Kooi, B. 2008. *Dynamic Epistemic Logic*, volume 337. Springer Coll. Synthese Library: Studies in Epistemology, Logic, Methodology, and Philosophy of Science, Dordrecht.

Veltman, F. 1996. Defaults in update semantics. *Journal of Philosophical Logic* 25:221–261.