Social Consolidations

David Santos, Yuri

Published in:
Foundations of Information and Knowledge Systems

DOI:
10.1007/978-3-030-39951-1_4

Social Consolidations: Rational Belief in a Many-Valued Logic of Evidence and Peerhood

Yuri David Santos(✉)

University of Groningen, Groningen, The Netherlands
y.david.santos@rug.nl

Abstract. We explore an interpretation of FVEL, a four-valued logic of evidence, where states represent agents, the propositional layer corresponds to the evidence available to these agents, and the relation corresponds to peerhood connections between them. Belief is determined based on the agent's evidence, but also on her peers' evidence. *Consolidation* functions are proposed, which map evidence situations to belief attitudes. We adapt some postulates of Social Choice Theory to our belief formation setting and, with them, we separate rational from irrational consolidations. We define a dynamic operator for addition and removal of evidence, which serves as a basis for some essential dynamic postulates and also for future developments on consolidations that take amounts of evidence into account. Our main technical result is a characterisation of a class of consolidations satisfying most of our rationality postulates.

Keywords: Evidence logics · Epistemic logic · Many-valued logic

1 Introduction

Four-valued epistemic logic (FVEL) [29] was first designed to model scenarios where agents are uncertain about the evidence publicly available. Here we give another interpretation to this logic, where the binary relation represents peerhood connections. Therefore, each state will represent the evidential state of one agent. This puts this work in line with other network logics such as [4,7].

In our setting, agents have four-valued evidence for propositions, embodied by a four-valued valuation function over atoms, which represents only evidence *for* that atom, only evidence *against* it, evidence both *for and against* it, or no evidence at all. Our main goal in this paper is to find rational ways of forming beliefs for these agents, given their own evidence and their peers'. With that in mind, we establish some rationality postulates and check some definitions of belief that respect those postulates, and some that do not.

A. Herzig and J. Kontinen (Eds.): FoIKS 2020, LNCS 12012, pp. 58–78, 2020.
https://doi.org/10.1007/978-3-030-39951-1_4

After that, we introduce a dynamic operator for addition/removal of evidence. This operator is used to axiomatise some of the postulates, but also to define two new ones, which serve to rule out some undesirable consolidations. We then prove that these axioms characterise a class of consolidations satisfying most of the main postulates. Finally, we show how this operator can be used to "count" peers, which in the future can be employed to define consolidations that form beliefs based on the amount of evidence for or against something.[1]

2 Syntax and Semantics

In this section we explore a variant of *four-valued epistemic logic* (FVEL) [29].

2.1 Syntax

Let At be a countable set of atoms. Below, $p \in At$; the classical part of the language, $\mathscr{L}_0$, is represented below by ψ; the propositional part $\mathscr{L}_1$, where $\psi \in \mathscr{L}_0$, is represented by χ; and the complete language $\mathscr{L}$, where $\psi \in \mathscr{L}_0$ and $\chi \in \mathscr{L}_1$, is given by φ:

$$\psi ::= p \mid {\sim}\psi \mid (\psi \wedge \psi) \qquad \chi ::= \psi \mid \neg\chi \mid (\chi \wedge \chi) \mid {\sim}\chi$$

$$\varphi ::= \chi \mid {\sim}\varphi \mid (\varphi \wedge \varphi) \mid \Box\varphi \mid B\psi$$

We abbreviate $\varphi \vee \psi \stackrel{\text{def}}{=} {\sim}({\sim}\varphi \wedge {\sim}\psi)$ and $\Diamond\varphi \stackrel{\text{def}}{=} {\sim}\Box{\sim}\varphi$. We restrict belief to classical propositional formulas ($\mathscr{L}_0$) because formulas with $\neg$ refer to evidence, and we do not want agents forming beliefs about evidence, only about facts.

Formulas such as p are read as *the agent has evidence for* p, whereas $\neg p$ is read as *the agent has evidence against* p, and ${\sim}\varphi$ as *it is not the case that* φ. We read $\Box\varphi$ as φ *holds for all peers* and $B\varphi$ as *the agent believes* φ.[2, 3]

2.2 Semantics

Models are tuples $M = (S, R, V)$, where S is a finite set of agents, R is a binary relation on S representing "peerhood" and $V : At \times S \to \mathscr{P}(\{0, 1\})$ is a four-valued valuation representing agents' evidence: $\{1\}$ is *true* (t), $\{0\}$ is *false* (f),

[1] Some proofs at: https://github.com/ydsantos/appendix_scons/blob/master/proofs.pdf.

[2] Notice that our language is non-standard in that even though a formula in $\mathscr{L}_1$ has an evidential meaning (such as p meaning *the agent has evidence for* p), under the belief operator B these formulas are read as factual statements (e.g. Bp means that *the agent believes* p and not that *the agent believes that she has evidence for* p).

[3] We chose B (belief) instead of K (knowledge) because we are working with imperfect evidence, which can be misleading. Therefore, our agents can form false beliefs, which violate factivity, a standard requirement for knowledge.

$\{0,1\}$ is *both* (b) and $\emptyset$ is *none* (n). A satisfaction relation is defined as follows:

$$M, s \models p \text{ iff } 1 \in V(p,s) \qquad M, s \models \neg p \text{ iff } 0 \in V(p,s)$$
$$M, s \models {\sim}\varphi \quad \text{iff } M, s \not\models \varphi$$
$$M, s \models (\varphi \wedge \psi) \quad \text{iff } M, s \models \varphi \text{ and } M, s \models \psi$$
$$M, s \models \neg(\varphi \wedge \psi) \quad \text{iff } M, s \models \neg\varphi \text{ or } M, s \models \neg\psi$$
$$M, s \models \Box\varphi \quad \text{iff for all } t \in S \text{ s.t. } sRt, \text{ it holds that } M, t \models \varphi$$
$$M, s \models \neg{\sim}\varphi \text{ iff } M, s \models \varphi \qquad M, s \models \neg\neg\varphi \text{ iff } M, s \models \varphi$$

Note that the semantics of $\neg$ is defined in a case-by-case fashion (this operator comes from FDE [5]). An extended valuation function $\overline{V}$ can be defined differently for each type of formula. If $\varphi \in \mathscr{L}_1$, then: $1 \in \overline{V}(\varphi, s)$ iff $M, s \models \varphi$; $0 \in \overline{V}(\varphi, s)$ iff $M, s \models \neg\varphi$. Otherwise: $1 \in \overline{V}(\varphi, s)$ iff $M, s \models \varphi$ iff $0 \notin \overline{V}(\varphi, s)$.

As pointed out in [29], this logic can be seen as a modal extension of FDE [5], with the addition of a classical negation. The logic FDE deals with evidence differently than other logics such as intuitionistic logic [18,33]. While both are weaker than classical logic, the concept of justification as existence of constructive proofs is much stronger than what we consider evidence in this paper. In our case, evidence can be misleading, as mentioned before. FDE is more suitable for modelling situations with incomplete and inconsistent evidence, while FVEL extends this logic to a modal setting, enabling us to talk about multiple agents. FVEL also includes a classical negation, which gives it much more expressive power, and many of the definitions and results in this paper make use of this operator (${\sim}$). Among other things, it allows us to define formulas discriminating which of the four truth values a formula $\varphi \in \mathscr{L}_1$ has: $\varphi^n \stackrel{\text{def}}{=} ({\sim}\varphi \wedge {\sim}\neg\varphi)$; $\varphi^f \stackrel{\text{def}}{=} {\sim}{\sim}({\sim}\varphi \wedge \neg\varphi)$; $\varphi^t \stackrel{\text{def}}{=} {\sim}{\sim}(\varphi \wedge {\sim}\neg\varphi)$; $\varphi^b \stackrel{\text{def}}{=} {\sim}{\sim}(\varphi \wedge \neg\varphi)$. In words, a formula φ^x is satisfied ($M, s \models \varphi^x$) iff φ has value $x \in \{t, f, b, n\}$, i.e. $\overline{V}(\varphi, s) = x$.

We say that $\Sigma \models \varphi$ (*Σ entails φ*) when for all models M and states s, if $M, s \models \sigma$ for all $\sigma \in \Sigma$, then $M, s \models \varphi$. We say that $M \models \varphi$ if $M, s \models \varphi$ for all states s of M. And $\models \varphi$ (φ is *valid*) if $M \models \varphi$ for all M; otherwise φ is *invalid*. If $\models {\sim}\varphi$, we say φ is *contradictory*, and if φ is not contradictory nor valid, it is *contingent*. If a formula is valid or contingent, it is *satisfiable*. Call the *truth range* of φ the set $\{x \mid$ there is a model $M = (S, R, V)$ and an $s \in S$ s.t. $\overline{V}(\varphi, s) = x\}$. The following result will be useful for some of the proofs:

Proposition 1. *All formulas in $\mathscr{L}_0$ have one of the following four truth ranges: $\{\{1\}\}$, $\{\{0\}\}$, $\{\{0\}, \{1\}\}$, $\{\emptyset, \{0\}, \{1\}, \{0,1\}\}$. A formula in $\mathscr{L}_1$ can have any truth range in $\mathscr{P}(\mathscr{P}(\{0,1\})) \setminus \emptyset$ except for $\{\emptyset\}$, $\{\{0,1\}\}$, and $\{\emptyset, \{0,1\}\}$.*

The central question of this paper is how to define the semantics for belief based on the evidence, a process we call *consolidation* (see [30]). A key philosophical assumption of this project is that *rational belief is determined by evidence.*

3 Rationality Conditions for Consolidations

3.1 Epistemic Autonomy Versus Epistemic Authority

In social epistemology, there is currently a lot of debate around the topics of peer disagreement and higher-order evidence [6,12,13,17,22–24]. One important question in this debate is: *What should a rational agent do when her peers – who she deems as rational as her – have different opinions on some proposition?* There are many different proposals in the literature as to what to do in this case. Nevertheless, we can roughly categorise them into two main groups: the *equal weight* views [9], and the *steadfast* views [21]. The former tend to consider the agent and her peers to be on equal footing, so if you and your peer disagree on something, your opinion should be something in the middle of both opinions. The latter claim that you are entitled to trust yourself more than you trust your peers – maybe because you have direct access to your evidence, as opposed to mere testimonial access to your peers' evidence, or because of some other reason. In both views, the concept of *peerhood* is preeminent. It is assumed that, in what matters, you and your peers are of equal competence. Evidently, if one's peer is far more competent than oneself in the topic at hand and one knows that, the rational thing to do is to defer to her judgement (but in that case she is *not* your peer). What enables peerhood is the lack of such higher-order knowledge: we usually do not know exactly how competent a peer is, so the reasonable (and modest) thing to do is to assume that the relevant people in the given case are (possibly) as competent as you, except if you have a "defeater" for that belief[4].

3.2 Rationality Postulates

Now we propose and discuss a series of rationality postulates, mostly adapted from postulates from Social Choice Theory (SCT) [2,14,31]. SCT is concerned with determining outcomes of voting from certain voting profiles. The adaptation we make here is in the sense that a rational belief in propositions (atomic or otherwise) will be determined from the evidence possessed by the agent and her peers, so here "voting profiles" become evidence, and "election outcome" becomes belief attitude. Consolidations are not voting procedures, but involve the weighing of inputs to find a suitable outcome.[5]

Regardless of the semantics of B, which is not yet defined, the following function Att serves as a shorthand for the doxastic attitude of an agent s w.r.t. a formula φ (belief, disbelief or abstention):

Definition 1 (Attitude). *Let* $\mathrm{Att} : \mathscr{L}_0 \times S \rightarrow \{1, 0, -1\}$ *be a function such that:* $\mathrm{Att}(\varphi, s) = 1$ *iff* $M, s \models B\varphi$*;* $\mathrm{Att}(\varphi, s) = -1$ *iff* $M, s \models B{\sim}\varphi$*; otherwise* $\mathrm{Att}(\varphi, s) = 0$.[6]

[4] As a scientist investigating hypothesis H, you consider another scientist also investigating H to be your peer, but not if she committed fraud in the past.

[5] Note, however, that we only make a loose connection to SCT here, not a formal one.

[6] The function Att also depends on a model M, but this will be left implicit. We will usually write Att' if we are referring to another model M', Att'' for M'', and so on.

Postulate 1 (Consistency (Con)). *For all models M and $s \in S$: let $\Sigma = \{\varphi \in \mathscr{L}_0 \mid M, s \models B\varphi\}$. Then $\Sigma \not\models p \wedge {\sim} p$.*

The postulate above is the most important demand on our consolidations: rational belief has to be consistent.

Postulate 2 (Modesty (Mod)). *For all models $M = (S, R, V)$, all $s \in S$, and all contingent $\varphi \in \mathscr{L}_0$, there is a model $M' = (S', R', V')$ with $S \subseteq S'$ s.t. $\mathrm{Att}(\varphi, s) \neq \mathrm{Att}'(\varphi, s)$, where $V|_s = V'|_s$.*[7]

Postulate 2 says that it is possible to change an agent's attitude toward a contingent formula just by changing her peerhood connections and the evidence of her peers. Modesty is adapted from the SCT postulate of *non-dictatorship*: the outcome of the election is not determined by one single agent. Postulate 3 also comes from non-dictatorship, but for *Modesty* we think of the agent as her own dictator.

The plausibility of this postulate hinges on the plausibility of the claim that regardless of what evidence you have, it is never rational to ignore others' evidence. This, in turn, depends on the outcome of the debate in epistemology discussed above. In any case, is the format of this postulate adequate? The restriction to contingent formulas seems justified: if we reject *Logical Omniscience*, it might be acceptable to abstain from judgement on tautologies and contradictions, but it seems irrational to expect one to be persuaded to abandon a belief in a tautology or adhere to a contradiction. Keeping $V|_s$ untouched captures exactly the idea of not changing one's evidence, but possibly changing others'. The $S \subseteq S'$ part demands that the original agents be preserved. This is innocuous, for even if a change in belief demands the removal of a peer, that can be obtained by removing the connection (changing R); non-peers do not matter in our setting. A stronger variant of Modesty could be considered, *Strong Modesty*, where not only is it possible to change the attitude for any formula, but also any other attitude is possible. This could be plausible, but expecting a radical change in attitude (for example, from disbelief to belief) for any contingent proposition might require a huge *amount* of evidence, but we are not representing this *aspect* of evidence here; we do make a step in this direction in Sect. 5.

Postulate 3 (No Gurus (NG)). *For all agents $s, t \in S$ (with $s \neq t$) and all contingent $\varphi \in \mathscr{L}_0$, there is a model $M = (S, R, V)$ s.t. $\mathrm{Att}(\varphi, s) \neq \mathrm{Att}(\varphi, t)$.*

This postulate says that for any formula there is a model such that the attitudes of two agents towards that formula differ, i.e. an agent's opinion is not determined by anyone else's. This postulate also stems from the postulate of non-dictatorship in SCT (in a more obvious way). We have that a consolidation satisfying Mod also satisfies NG (see Proposition 4 later). So if Modesty is plausible, then this postulate has to be as well. In principle, it might be odd to think that, for example, two biologists could rationally disagree on whether natural selection

[7] We denote by $V|_s$ the restriction of a valuation V to $At \times \{s\}$, with $s \in S$.

happens. This apparent controversy is only superficial, though. If we stick to our key assumption that evidence determines rational belief, then that should be possible *given* they have access to different circles – with one of them possibly possessing misleading evidence.

Postulate 4 (Equal Weight (EW)). *Consider any model* $M = (S, R, V)$*, any two agents* $s, t \in S$*, and a valuation* V' *such that* $V'(p, s) = V(p, t)$*,* $V'(p, t) = V(p, s)$*, and* $V'(p, u) = V(p, u)$ *for all* $u \in S \setminus \{s, t\}$*, for all* $p \in At$*. Then if* sRt*, it holds that* $\mathrm{Att}'(\varphi, s) = \mathrm{Att}(\varphi, s)$*, for all* $\varphi \in \mathscr{L}_0$*.*

What this postulate says is that if you swap all your evidence with the evidence of one of your peers, your beliefs do not change: you treat your evidence and your peers' equally. It comes from the SCT postulate of *anonimity*: if we have the same voting profile but swap the voters, the outcome does not change. Again, the plausibility of this postulate depends on your position in the debate of Sect. 3.1.

Postulate 5 (Atom Independence (AI)). *Consider any model* $M = (S, R, V)$*. For any atom* $p \in At$*, if* V' *is a valuation s.t.* $V'(p, s) = V(p, s)$ *for all* $s \in S$*, then* $\mathrm{Att}(p, s) = \mathrm{Att}'(p, s)$ *for all* $s \in S$*.*

The valuation of one atom should not interfere in the attitudes towards another. This postulate is adapted from *independence of irrelevant alternatives*: the outcome between x and y should only depend on voters opinions w.r.t. x and y; changing the preferences between other candidates does not affect the outcome. A more "local" version of this postulate could be formulated: for any $p \in At$ and $s \in S$, if $V(p, s) = V'(p, s)$ and $V(p, t) = V'(p, t)$ for all t such that sRt, then $\mathrm{Att}(p, s) = \mathrm{Att}'(p, s)$.[8] We can prove that this definition is equivalent to AI.

Let $\preceq$ be the smallest reflexive and transitive relation $\preceq$: $\mathscr{P}(\{0, 1\}) \times \mathscr{P}(\{0, 1\})$ such that $\{0\} \preceq \emptyset$, $\{0\} \preceq \{0, 1\}$, $\emptyset \preceq \{1\}$ and $\{0, 1\} \preceq \{1\}$. Let $\npreceq$ be the complement of $\preceq$, and define $x \prec y$ iff $x \preceq y$ and $y \npreceq x$.

Postulate 6 (Monotonicity (Mon)). *Consider a model* $M = (S, R, V)$ *and a* V' *which coincides with* V*, except that* $V'(p, s) \neq V(p, s)$ *for one* $s \in S$ *and* $p \in At$*. If* $V(p, s) \prec V'(p, s)$*, then for all* $t \in S$*,* $\mathrm{Att}(p, t) \leq \mathrm{Att}'(p, t)$*. If* $V'(p, s) \prec V(p, s)$*, then for all* $t \in S$*,* $\mathrm{Att}'(p, t) \leq \mathrm{Att}(p, t)$*.*

Postulate 6 states that if the valuation only changes positively/negatively for one atom and one agent, then the attitude towards this atom for any agent should either stay the same, or change according to the same trend (more positive/negative). Monotonicity was adapted from a homonymous SCT postulate: if a profile is altered only by promoting (demoting) one candidate, the outcome should either change only by promoting (demoting) this candidate, or not change.

[8] We thank an anonymous reviewer for this suggestion.

Now there is a question of adequacy of the format of this postulate. There is not always a unique way of changing a valuation to produce a certain change in the (extended) valuation of a complex formula, so we limited this postulate to atomic changes. The other question regarding format is why the postulate limits the valuation change to only one atom and one agent. Clearly changing one atom in one direction (according to $\prec$) for more agents, or changing several atoms in this fashion, should preserve monotonicity. These "cumulative" effects are already covered by the postulate as it is.

Postulate 7 (Doxastic Freedom (DF)). *Consider any set of agents S and any function $f : At \times S \to \{1, -1, 0\}$. Then there is a model $M = (S, R, V)$ such that* $\mathrm{Att}(p, s) = f(p, s)$ *for all $p \in At$ and $s \in S$.*

DF says that any combination of attitudes towards atoms is possible for any agent. It is adapted from *non-imposition*: every outcome is achievable by some voting profile. This postulate seems somehow connected to AI. However:

Observation 1. *A consolidation satisfying Doxastic Freedom does not necessarily satisfy Atom Independence. The converse also holds.*

Postulate 8 (Consensus (Css)). *If for some agent $s \in S$ and some $\varphi \in \mathscr{L}_0$ we have that $\overline{V}(\varphi, s) = \{1\}$ (or $\{0\}$), and for all $t \in S$ such that sRt: $\overline{V}(\varphi, t) = \{1\}$ (or $\{0\}$), then* $\mathrm{Att}(\varphi, s) \neq -1$ *(or 1).*

Consensus is derived from the SCT postulate of *unanimity*: if all voters prefer one candidate over another, then so must the outcome. It says that if an agent and all her peers have unambiguous evidence about some atom, then she should not believe contrary to that. We can define *Strong Consensus* in a similar way, but instead of demanding no contrary belief, it demands belief in case of unanimous positive evidence and disbelief in case of unanimous negative evidence.

Observation 2. *A consolidation satisfying Strong Consensus and Consistency also satisfies Consensus.*

Proof. One just has to see that $M, s \models B\varphi$ implies $M, s \not\models B{\sim}\varphi$ for a consolidations satisfying Consistency (and similarly for the $B{\sim}\varphi$ case). □

Notice that this stronger variant, in combination with Proposition 1, entails a form of *logical omniscience*. We could also have defined the postulate differently by considering unanimity among all agents instead of one agent and her peers, but, again, we are assuming that non-peers are inaccessible/irrelevant.

Postulate 9 (Logical Omniscience (LO)). *For all models M and $s \in S$: if $\Sigma \models \varphi$ and $M, s \models B\sigma$ for all $\sigma \in \Sigma$, then $M, s \models B\varphi$.*

This postulate is not derived from any postulate of SCT. It is debatable whether it should be satisfied or not, but as a normative demand on real agents we consider it too strong. Notice that it implies the knowledge of all validities, as they are consequences of the empty set, and also that the doxastic state has to be consistent or it will be trivialised.

In summary, all the postulates listed in this section are expected to be satisfied by any rational consolidation (call these *core postulates*), except for Mod and EW, whose normative status depend on the reader's philosophical commitments w.r.t. the debate of Sect. 3.1, and LO, which is also part of another long debate [11,19,20,26]. No impossibility theorem à la Arrow [2] ensues, and consolidations satisfying all core postulates are presented. One main difference of our approach that might explain this is that we do not have preference orders over attitudes. Note also that our connection to SCT is not fully formal, our postulates are only inspired by it.

4 Social Consolidations

In this section we will define *consolidation policies*, that is, methods of defining belief from evidence. We expect the most reasonable consolidations to satisfy all the core postulates, and unreasonable ones to violate at least one of them.

4.1 Preliminaries

Before talking about consolidations, we will formally specify what are the possible ones. Now let $\mathbb{M} = \{(M, s) \mid M = (S, R, V)$ is an FVEL model and $s \in S\}$ be the class of all *pointed models*. First, we draw the following definition from the literature on n-bisimulations:

Definition 2 (1-Bisimulation). *Consider two FVEL models $M = (S, R, V)$ and $M' = (S', R', V')$, an $s \in S$ and an $s' \in S'$. We say that $(M, s) \leftrightharpoons (M', s')$, read (M, s) is* 1-bisimilar *to (M', s'), iff:*

atoms *For all $p \in At$, $V(p, s) = V'(p, s')$;*

back *For all $t' \in S'$ s.t. $s'R't'$, there is a $t \in S$ s.t. sRt and $V(p, t) = V'(p, t')$ for all $p \in At$.*

forth *For all $t \in S$ s.t. sRt, there is a $t' \in S'$ s.t. $s'R't'$ and $V(p, t) = V'(p, t')$ for all $p \in At$.*

The purpose of Definition 2 is to determine whether two pointed models have equivalent evidence. Since our relation R of peerhood is not transitive, we assume that our agents only have access to their own evidence and their peers'. So formulas such as $\Box p$ are relevant for consolidation, whereas $\Box\Box p$ is not.

Proposition 2. *$(M, s) \leftrightharpoons (M', s')$ implies: $M, s \models \varphi$ iff $M', s' \models \varphi$ for all $\varphi \in \mathscr{L}$ not containing B nor nested $\Box$. The converse also holds for* image-finite *models (each agent has finitely many peers).*

Proposition 3. *The relation $\leftrightharpoons$ is an equivalence relation.*

Then $\leftrightharpoons \subseteq \mathbb{M} \times \mathbb{M}$. Denote by $[M, s]$ the equivalence class of (M, s) under $\leftrightharpoons$, that is, $[M, s] = \{(M', s') \in \mathbb{M} \mid (M, s) \leftrightharpoons (M', s')\}$. Let $\mathbb{M}/_{\leftrightharpoons}$ be the quotient class of $\mathbb{M}$ by $\leftrightharpoons$, that is, the class of equivalence classes of $\mathbb{M}$ under $\leftrightharpoons$. Then, we are interested in the following:

Definition 3. *A consolidation is a function* $\mathbb{C} : \mathbb{M}/_{\leftrightharpoons} \times \mathscr{L}_0 \to \{0, 1\}$. *For any model* $M = (S, R, V)$ *with* $s \in S$, *we set* $M, s \models B\varphi$ *iff* $\mathbb{C}([M, s], \varphi) = 1$.

Proposition 4. *A consolidation satisfying Mod also satisfies NG.*

With these definitions in hand, we will introduce the following:

Definition 4. *We say that a condition is* axiomatisable *when: it holds iff all* $\sigma \in \Sigma$ *are valid, for some* $\Sigma \subseteq \mathscr{L}$. *We say that a condition is* negatively axiomatisable *when: it holds iff all* $\sigma \in \Sigma$ *are invalid, for some* $\Sigma \subseteq \mathscr{L}$.

Proposition 5. *Consistency holds iff for all finite* $\Sigma = \{\sigma_1, ...\sigma_n\} \subseteq \mathscr{L}_0$ *such that* $\Sigma \models p \wedge {\sim}p$, ${\sim}(B\sigma_1 \wedge ... \wedge B\sigma_n)$ *is valid.*

Proof. The logic of $\mathscr{L}_0$ is basically classical propositional logic (as mentioned in [29]), and is, therefore, compact. So for any $\Sigma \models \varphi$ with $\varphi \in \mathscr{L}_0$, there is a finite $\Sigma' \subseteq \Sigma$ such that $\Sigma' \models \varphi$. The case where $\varphi = p \wedge {\sim}p$ is a particular case of this. So all inconsistent subsets of $\mathscr{L}_0$ have a finite inconsistent subset. □

Proposition 6. *Logical Omniscience holds iff for all finite* $\Sigma = \{\sigma_1, ...\sigma_n\} \subseteq \mathscr{L}_0$ *and* $\varphi \in \mathscr{L}_0$ *such that* $\Sigma \models \varphi$, ${\sim}(B\sigma_1 \wedge ... \wedge B\sigma_n \wedge {\sim}B\varphi)$ *is valid.*

Proof. The reasoning is similar to the case for Proposition 5. □

Note that Propositions 5 and 6 follow from compactness of $\mathscr{L}_0$. Now consider the following axioms:

$$\textbf{C1} \quad {\sim}((\varphi^t \wedge \Box\varphi^t) \wedge B{\sim}\varphi) \qquad \textbf{C2} \quad {\sim}((\varphi^f \wedge \Box\varphi^f) \wedge B\varphi)$$

Proposition 7. *A consolidation satisfying Consistency satisfies Consensus iff* ***C1*** *and* ***C2*** *are valid.*

4.2 Consolidation Policies

First, we will look at the most straightforward (and naive) possibility: $M, s \models B\varphi$ iff $M, s \models \Box\varphi$. This possibility is appealing because it is familiar and simple. First, let us note that, in order to include the evidence of the agent itself in the consolidation, we have to require the model to be reflexive. This raises the question: is the agent a peer of herself (see [9])? If yes, then we should only work with reflexive models, if not, then only with *anti-reflexive* models (sRs holds for no s). This is not so crucial as we can (and will) use an equivalent definition for anti-reflexive models: $M, s \models B\varphi$ iff $M, s \models \varphi \wedge \Box\varphi$. So we assume that agents are not peers of themselves. We call this latter definition *naive consolidation*.

Proposition 8. *Naive consolidation satisfies Con, Mod, EW, AI, Mon and Strong Css. It does not satisfy DF and LO.*

Surprisingly, naive consolidation only fails one core postulate: Doxastic Freedom. It is surprising because this consolidation actually *ignores all negative evidence.*

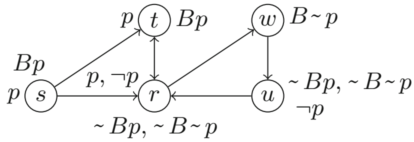

Fig. 1. An example of naive consolidation. Agent s believes p, but not $\sim p$, since all her peers and herself satisfy p (have evidence for p), and not $\sim p$. One of the peers (r) has $\neg p$, but s ignores that. Agent w believes $\sim p$, even though she does not have evidence against p. She believes $\sim p$ only on the grounds that she and r do not have evidence for p. Agent u does not believe p nor $\sim p$, because she does not have evidence for p, but her only peer does.

Another simple consolidation we can analyse is the *sceptical consolidation*, which sets $M, s \not\models B\varphi$ for all $\varphi \in \mathscr{L}_0$. Fortunately this extreme position is blocked by two of our core postulates.

Proposition 9. *Sceptical consolidation satisfies Con, EW, AI, Mon and Css. It does not satisfy in general NG (and therefore Mod), DF and LO.*

Now we will try a more sophisticated definition:

Definition 5. *Call $\mathscr{C}$-consolidations the policies defined by:*

$$\begin{array}{ll} M, s \models Bp & \textit{iff } \mathscr{C}(V_p^s, V_{\neg p}^s, V_{\Diamond p}^s, V_{\Diamond \neg p}^s, V_{\Box p}^s, V_{\Box \neg p}^s) = 1 \\ M, s \models B\sim p & \textit{iff } \mathscr{C}(V_p^s, V_{\neg p}^s, V_{\Diamond p}^s, V_{\Diamond \neg p}^s, V_{\Box p}^s, V_{\Box \neg p}^s) = -1 \\ M, s \models B\sim\sim\varphi & \textit{iff } M, s \models B\varphi \\ M, s \models B(\varphi \wedge \psi) & \textit{iff } M, s \models B\varphi \textit{ and } M, s \models B\psi \\ M, s \models B\sim(\varphi \wedge \psi) & \textit{iff } M, s \models B\sim\varphi \textit{ or } M, s \models B\sim\psi \end{array}$$

where V_χ^t is 1 if $1 \in \overline{V}(\chi, t)$ and 0 otherwise; and $\mathscr{C} : \{0,1\}^6 \to \{1, -1, 0\}$ is a function that maps evidence (in this case represented by the six binary parameters) to a belief attitude (1 for belief, -1 for disbelief and 0 for abstention).

What is a good definition for $\mathscr{C}$? As we can see above, the real consolidation effort is only with respect to atomic propositions, while more complex beliefs are formed from those atomic beliefs. Some advantages of this approach are that it uses all evidence available for each atom, the agent still retains some inference power (with which it can derive other beliefs), and avoids malformed definitions, such as: $M, s \models B\varphi$ iff $M, s \models \varphi^t \wedge \Box\varphi^t$; $M, s \models B{\sim}\varphi$ iff $M, s \models \varphi^f \wedge \Box\varphi^f$. In words: the agent believes a formula if she and her peers have only positive evidence for it, and believes its negation if she and her peers have only negative evidence for it. This seems like a good (if too cautious) definition at first sight, but it is actually not well-formed. We can verify whether $B{\sim}\psi$ via the second clause, but also via the first if $\varphi = {\sim}\psi$. And these can sometimes give conflicting results. We avoid that by using $\mathscr{C}$ only to decide belief for literals. Moreover:

Proposition 10. *All $\mathscr{C}$-consolidations satisfy Con and AI.*

Our agents under $\mathscr{C}$-consolidations are not necessarily omniscient, but they present some properties related to unbounded logical power:

Proposition 11. *Consider any $\mathscr{C}$-consolidation, and a maximally consistent set of literals Σ. If $M, s \models B\sigma$ for all $\sigma \in \Sigma$ and $\Sigma \models \varphi$, then $M, s \models B\varphi$.*

Corollary 1. *Any $\mathscr{C}$-consolidation satisfying DF also satisfies NG.*

Proposition 12. *Belief in $\mathscr{C}$-consolidations is closed under* modus ponens*: if $M, s \models B\varphi$ and $M, s \models B{\sim}(\varphi \wedge {\sim}\psi)$, then $M, s \models B\psi$.*

Corollary 2. *Any $\mathscr{C}$-consolidation satisfies Logical Omniscience if we add the following clause to the semantics: if $\models \varphi$, then $M, s \models B\varphi$ (where $\varphi \in \mathscr{L}_0$).*

There are $3^{(2^6)} = 3^{64} \approx 3.43 \times 10^{30}$ consolidation function candidates for $\mathscr{C}$. The combinations $(0, 1, 1)$ for $V^s_{\Diamond p}, V^s_{\Diamond\neg p}, V^s_{\Box p}$ and $(1, 0, 1)$ for $V^s_{\Diamond p}, V^s_{\Diamond\neg p}, V^s_{\Box\neg p}$ are impossible, though, which leaves us with "only" $3^{48} \approx 7.98 \times 10^{22}$ relevantly different candidates. Now we consider some promising possibilities.

Policy I. Our first social consolidation policy is in Fig. 2. In cases of unambiguous evidence, the agent decides for belief or disbelief, accordingly. In the case of conflicting evidence, the agent already has some evidence, and since we want a consistent doxastic state, this entails that the agent will inevitably have to discard some evidence. So, in this case, the mere existence of evidence of one kind from one peer is enough to produce belief. However, when the agent has no evidence at all, even if she decides to abstain there is no waste of evidence, so she will be more demanding to change her view. In this case, unanimity of her peers is needed (see an example in Fig. 3).

Policy II. One might consider that our previous policy still does not justify the different treatment for the problematic evidence cases, and is therefore arbitrary. Hence, we can consider a second policy where the behaviour when the evidence is *none* imitates the case for *both*: consider a decision tree identical to that of Fig. 2 but with the subtree for none (the leftmost subtree) just replaced by that used for both (the rightmost one).

Proposition 13. *Policy I and II satisfy Monotonicity, Doxastic Freedom and Consensus. Modesty and Equal Weight are not satisfied.*

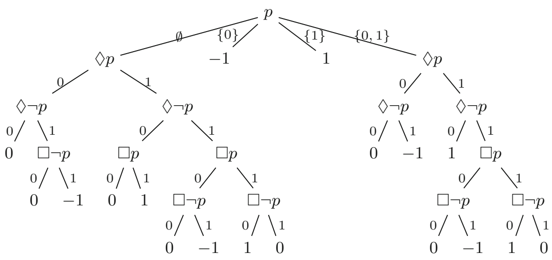

Fig. 2. Decision trees will be used to represent $\mathscr{C}$-consolidations. This one represents $\mathscr{C}$ for Policy I. Nodes are labelled by expressions that are representable with the six parameters for $\mathscr{C}$. The leaves are the outcomes of the consolidation: 1 for belief, −1 for disbelief and 0 for abstention of judgement.

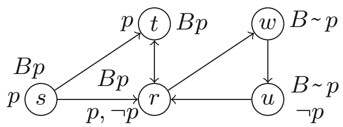

Fig. 3. Policy I applied to the model of Fig. 1. Here all agents except for r and w have unambiguous evidence about p, so they can easily form beliefs without looking at their peers. Agent w has no evidence whatsoever, so by the tree of Fig. 2 she decides to believe $\sim p$ due to her only peer satisfying $\neg p$. Agent r has evidence both for and against p. Since she has a peer with evidence for p, but no peer with evidence against p, she believes p. Note that by Fig. 2 this decision would have been different if r had no evidence at all.

Policy III. The previous policies are in the "steadfast" category. Our agent gives more weight to her own evidence than to others' opinions. We can devise a policy that is more in line with the "equal weight" view. In this case, we consider the relation R to be reflexive, and then "dissolve" the agents' exceptionality in the modal expression. Starting from the consolidation of Fig. 2, we can take its subtree for *both* as the decision tree for this policy (Fig. 4 (left)), ignoring the

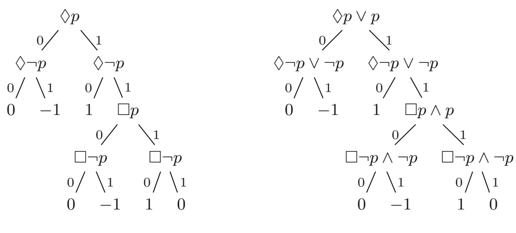

Fig. 4. Decision trees for $\mathscr{C}$ of Policy III for reflexive (left) and anti-reflexive (right) models. Both yield the same beliefs in their respective class of models.

inputs V^s_φ, $V^s_{\neg\varphi}$. This definition makes no distinction between the agent's own evidence and her peers'. We will, however, use the definition of Fig. 4 (right) instead, as we are working with anti-reflexive models. For an example of Policy III, see Fig. 5.

Proposition 14. *Policy III satisfies Mod, EW, Mon, DF and Css.*

5 Dynamics

The dynamic operations we will study use the following models for semantics:

Definition 6. *Consider a model $M = (S, R, V)$. We denote by $M^+_p = (S, R, V')$ any model s.t. for some $t \in S$, $V'(p, t) = V(p, t) \cup \{1\}$, and $V'(q, r) = V(q, r)$ when $q \neq p$ or $r \neq t$. We define M^-_p, $M^+_{\neg p}$, $M^-_{\neg p}$ analogously, but with $V'(p, t) = V(p, t) \setminus \{1\}$, $V'(p, t) = V(p, t) \cup \{0\}$, $V'(p, t) = V(p, t) \setminus \{0\}$, respectively.*

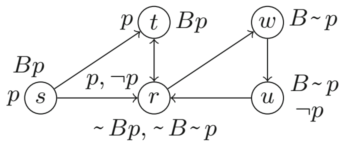

Fig. 5. Policy III applied to the model of Fig. 1. Agent w believes ~p because she or some peer have $\neg p$, but neither she nor her peer have p. All the other agents have evidence for and against p, either by themselves or via some peer. In this case, if the agent and all her peers have one type of evidence but not the other, a belief is formed. E.g. agent s and her peers have evidence for p but not all of them have $\neg p$, so she settles with belief in p. Agent r, on the other hand, has evidence for and against p (by herself or via a peer), but they are not unanimous about neither, therefore r abstains.

Now, with $l \in \{p, \neg p\}$ for some $p \in At$ and $\circ \in \{+, -\}$, we can define the following operator (with obvious additions to the language).

$$M, s \models [\circ l]\varphi \text{ iff for every model } M^\circ_l \text{ it holds that } M^\circ_l, s \models \varphi$$

So, for example, $M, s \models [+p]\varphi$ can be read as *if evidence for p is added for any agent, φ is the case for s.* A corresponding existential version of this operator can be defined by $\langle\circ l\rangle\varphi \overset{\text{def}}{=} {\sim}[\circ l]{\sim}\varphi$, with the expected semantics:

$$M, s \models \langle\circ l\rangle\varphi \text{ iff for some model } M^\circ_l \text{ it holds that } M^\circ_l, s \models \varphi$$

We note the following interactions between modalities:

$$M, s \models \Box[\circ l]\varphi \text{ iff } M, s \models [\circ l]\Box\varphi \qquad M, s \models \Diamond\langle\circ l\rangle\varphi \text{ iff } M, s \models \langle\circ l\rangle\Diamond\varphi$$

Interestingly, we can use the axioms below to define Monotonicity, revealing the hidden dynamic nature of that postulate.

M1	$\sim(Bp \wedge \langle +p \rangle \sim Bp)$	**M5**	$\sim(B{\sim}p \wedge \langle -p \rangle \sim B{\sim}p)$
M2	$\sim(Bp \wedge \langle -\neg p \rangle \sim Bp)$	**M6**	$\sim(B{\sim}p \wedge \langle +\neg p \rangle \sim B{\sim}p)$
M3	$\sim(\sim B{\sim}p \wedge \langle +p \rangle B{\sim}p)$	**M7**	$\sim(\sim Bp \wedge \langle -p \rangle Bp)$
M4	$\sim(\sim B{\sim}p \wedge \langle -\neg p \rangle B{\sim}p)$	**M8**	$\sim(\sim Bp \wedge \langle +\neg p \rangle Bp)$

Proposition 15. *A consolidation satisfying Consistency satisfies Monotonicity iff **M1**-**M8** are valid.*

Proof. ($\Leftarrow$) If $\sim(Bp \wedge \langle +p \rangle \sim Bp)$ is valid, then for any M, s, it holds that $M, s \not\models Bp$ or $M, s \not\models \langle +p \rangle \sim Bp$, which implies that $M, s \models Bp$ implies $M, s \not\models \langle +p \rangle \sim Bp$. This implies that if $M, s \models Bp$, then there is no M_p^+ such that $M_p^+, s \not\models Bp$. This covers one of the cases of Monotonicity. By analogous reasoning with the other axioms, we get all the other cases.

($\Rightarrow$) The axiom $\sim(Bp \wedge \langle +p \rangle \sim Bp)$ is valid if, for arbitrary M and s, $M, s \models Bp$ implies there is no M_p^+ such that $M_p^+, s \not\models Bp$. Indeed a model M_p^+ satisfies the condition $V(p,t) \preceq V'(p,t)$ for some t (by Definition 6). In this case Monotonicity implies that $\text{Att}'(p,s) \geq \text{Att}(p,s)$. So indeed, if $M, s \models Bp$, which by Consistency means that $\text{Att}(p,s) = 1$, we can only have $\text{Att}'(p,s) = 1$, so $M_p^+, s \models Bp$. So the semantic conditions for **M1** are satisfied. Notice that the case for **M2** is similar, because a model $M_{\neg p}^-$ also satisfies $V(p,t) \preceq V'(p,t)$ for some t. The cases for the other axioms are similar. □

We can do something similar for AI, where $l \in \{q, \neg q\}$ and $q \neq p$:

AI1	$\sim(Bp \wedge \langle \circ l \rangle \sim Bp)$	**AI3**	$\sim(\sim Bp \wedge \langle \circ l \rangle Bp)$
AI2	$\sim(B{\sim}p \wedge \langle \circ l \rangle \sim B{\sim}p)$	**AI4**	$\sim(\sim B{\sim}p \wedge \langle \circ l \rangle B{\sim}p)$

Proposition 16. *For image-finite models and a finite At, a consolidation satisfies Atom Independence iff **AI1**-**AI4** are valid. For infinite At, validity of **AI1**-**AI4** do not imply Atom Independence.*

Proof. ($\Leftarrow$) Suppose **AI1-AI4** are valid. If our models are image-finite and At is finite, then for any two models M and M', if there is a p such that for all $s \in S$ we have $V(p,s) = V'(p,s)$, then there is a finite sequence: $M, M_{l_1}^{\circ_1}, (M_{l_1}^{\circ_1})_{l_2}^{\circ_2}, ..., M'$, where $l_1, l_2, ...$ do not involve p. If $\text{Att}(p,s) \neq \text{Att}'(p,s)$ (for M and M', respectively), then there is one M_i in this sequence such that $\text{Att}_i(p,s) \neq \text{Att}_{i+1}(p,s)$. But if **AI1-AI4** are valid, this is not possible.

($\Rightarrow$) Assume that Atom Independence is satisfied, and $Bp \wedge \langle \circ l \rangle \sim Bp$ is satisfiable. Then there is a M_l° and s such that $M_l^\circ, s \not\models Bp$, while $M, s \models Bp$. But then $V_l^\circ(p,t) = V(p,t)$ for all t, but $\text{Att}(p,s) \neq \text{Att}_l^\circ(p,s)$, and therefore Atom Independence does not hold. Contradiction. Therefore **AI1** is valid. The other cases are similar.

Now we show a consolidation which satisfies **AI1-AI4** but violates Atom Independence (in a setting with infinite At). First, we will need to define some preliminary notions. Let M, s have a *p-canonical* valuation iff $V(p,s) = \{1\}$ and $V(p,t) = \{1\}$ for all t with sRt, and $V(q,s) = \{0\}$ and $V(q,t) = \{0\}$ for all t with sRt, for all $q \neq p$. The *p-canonical* model of M, s is a pointed model $M^\star, s$, where the valuation of $M^\star$ is such that $M^\star, s$ has a p-canonical valuation. For two pointed models M, s and M', s which differ only in V, define the *distance* between them to be the size of the sequence (similar to the one built in the first part of this proof) needed to go from M to M'. If no such sequence exists, the distance is infinite. We can easily show that (*) if $M, s \leftrightarrows M', s'$, then M, s is at a finite distance from its p-canonical model iff M', s' is at a finite distance from its p-canonical model. Now define a consolidation $\mathbb{C}$ as follows: $M, s \models Bp$ iff M, s is at a finite distance from its p-canonical model, and $M, s \not\models B\varphi$ for all non-atomic φ. This consolidation respects Definition 4, due to (*). Moreover, this definition violates Atom Independence, for if we take a p-canonical M, s (with $\text{Att}(p,s) = 1$) and change the valuation of infinitely many atoms (without changing p) to obtain M^*, s, this new pointed model is not at a finite distance from its p-canonical model M, s, and therefore $\text{Att}^*(p,s) \neq 1$. This violates Atom Independence. Axioms **AI1** to **AI4**, however, are valid. Suppose $M, s \models Bp$. Then M, s is at a finite distance from its p-canonical model. For $M, s \models \langle \circ l \rangle {\sim} Bp$ to be satisfied, there needs to be a M_l°, s such that $M_l^\circ, s \models {\sim} Bp$. But that would mean that M_l° is at an infinite distance from its p-canonical model. This is impossible, for M, s is p-canonical and M_l° only differs from it in one atom for one agent. □

The following formula means that *there is an agent other than myself such that if we add/remove evidence l for her, φ holds* (where $l \in \{p, \neg p\}$, for some $p \in At$):

$$\langle\langle \circ l \rangle\rangle \varphi \stackrel{\text{def}}{=} (p^t \wedge \langle \circ l \rangle (p^t \wedge \varphi)) \vee (p^f \wedge \langle \circ l \rangle (p^f \wedge \varphi)) \vee (p^b \wedge \langle \circ l \rangle (p^b \wedge \varphi)) \vee (p^n \wedge \langle \circ l \rangle (p^n \wedge \varphi))$$

The two following postulates could have been defined before, but now we can define them less cumbersomely:

ES1 $Bp \wedge \langle + \neg p \rangle {\sim} Bp$

ES2 $B{\sim}p \wedge \langle - \neg p \rangle {\sim} B{\sim}p$

ES3 $Bp \wedge \langle - p \rangle {\sim} Bp$

ES4 $B{\sim}p \wedge \langle + p \rangle {\sim} B{\sim}p$

SS1 $Bp \wedge \langle\langle + \neg p \rangle\rangle {\sim} Bp$

SS2 $B{\sim}p \wedge \langle\langle - \neg p \rangle\rangle {\sim} B{\sim}p$

SS3 $Bp \wedge \langle\langle - p \rangle\rangle {\sim} Bp$

SS4 $B{\sim}p \wedge \langle\langle + p \rangle\rangle {\sim} B{\sim}p$

Postulate 10 (Evidence Sensitivity (ES)). ***ES1-ES4*** *are satisfiable.*

Postulate 11 (Social Sensitivity (SS)). ***SS1-SS4*** *are satisfiable.*

Observation 3. *A consolidation satisfying SS also satisfies ES.*

Now from Propositions 5, 7, 15–16 and Postulate 10–11, we get our main technical result:

Corollary 3. *A consolidation satisfies Consistency, Monotonicity, Consensus, Evidence Sensitivity and Social Sensitivity iff:* $\sim(B\sigma_1 \wedge ... \wedge B\sigma_n)$ *is valid, for all finite* $\Sigma = \{\sigma_1, ...\sigma_n\} \subseteq \mathscr{L}_0$ *such that* $\Sigma \models p \wedge \sim p$*;* ***M1-M8***, ***C1-C2*** *are valid; and* ***ES1-ES4***, ***SS1-SS4*** *are satisfiable.*

Atom Independence can be included (with its respective axioms **AI1-AI4**) if we apply the restrictions of Proposition 16. The significance of Corollary 3 is that it characterises a class of consolidations satisfying almost all core postulates. We conjecture that Doxastic Freedom and No Gurus are not axiomatisable (nor negatively so). A hint of why that might be the case for NG is that it is equivalent to saying that there is a model s.t.: $(M, s \models B\varphi$ and $M, t \not\models B\varphi)$ or $(M, s \models B\sim\varphi$ and $M, t \not\models B\sim\varphi)$ or $(M, s \not\models B\varphi$ and $M, t \models B\varphi)$ or $(M, s \not\models B\sim\varphi$ and $M, t \models B\sim\varphi)$. Our language, however, can only talk of belief from an agent's perspective, or modally (e.g. $\Diamond B\varphi$ – *there is a peer who believes* φ).

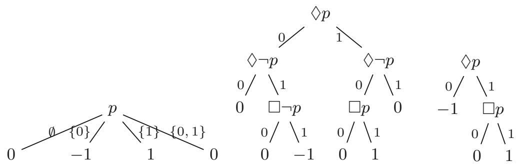

Fig. 6. Anti-social consolidation (left), Policy IV (center), and Policy V (right).

Figure 6 defines three more $\mathscr{C}$-consolidations which will show the importance of the new postulates. First, Social Sensitivity is the only core postulate to rule out anti-social consolidation, an unacceptable function that only takes the agent's own evidence into account.

Proposition 17. *Anti-social consolidation satisfies Mon, DF, Css and ES. Mod, EW and SS are not satisfied.*

Now it can be speculated that Evidence Sensitivity can be forced by a combination of other postulates, such as Strong Modesty, Atom Independence and Monotonicity. Policy IV satisfies all those postulates:

Proposition 18. *Policy IV satisfies Strong Modesty, Monotonicity, Doxastic Freedom, Consensus and Social Sensitivity. It does not satisfy Equal Weight.*

But that logical connection between those postulates does not hold. Interestingly, Policy IV violates Equal Weight, but this time not by the agents not giving enough importance to their peers, but by failing to appreciate their own evidence.

Policy V, which is just a modified version of naive consolidation that satisfies Doxastic Freedom, violates Evidence Sensitivity, because, as its cousin, it completely ignores negative evidence. What Evidence Sensitivity enforces is exactly this: that all evidence is taken into account at least in some occasions.

Proposition 19. *Policy V satisfies Strong Mod, Mon, DF and Css. It does not satisfy EW and ES.*

Proposition 20. *Policies I, II and III satisfy Social Sensitivity. Naive and sceptical consolidations do not satisfy Evidence Sensitivity.*

A summary of the consolidations appears in Table 1. But the main conclusion is that indeed the straightforward definitions such as naive and sceptical consolidations are very unsatisfactory, and the best ones (the only ones satisfying all core postulates) are Policies I-IV, depending on whether one adheres to equal weight or steadfast views.

Table 1. Postulates satisfied by consolidations. A.-S. is anti-social consolidation.

	Naive	Scept.	A.-S.	Pol. I	Pol. II	Pol. III	Pol. IV	Pol. V
AI	✓	✓	✓	✓	✓	✓	✓	✓
Mon	✓	✓	✓	✓	✓	✓	✓	✓
Css	✓	✓	✓	✓	✓	✓	✓	✓
NG	✓		✓	✓	✓	✓	✓	✓
DF			✓	✓	✓	✓	✓	✓
ES			✓	✓	✓	✓	✓	
SS				✓	✓	✓	✓	
Mod	✓					✓	✓	✓
EW	✓	✓				✓		

The $[\circ l]$ operators make the language more expressive, so we cannot use reduction axioms to obtain equivalent non-dynamic formulas. With these operators we gain the power to *count peers*[9]. Let us abbreviate $\langle \circ l \rangle ... \langle \circ l \rangle$, repeated n times, by $\langle \circ l \rangle^n$, with $\langle \circ l \rangle^0 \varphi \overset{\text{def}}{=} \varphi$. Then, with $l \in \{p, \neg p\}$ for some $p \in At$, we have:

$M, s \models {\sim} \langle -l \rangle^n \Box {\sim} l$ iff s has more than n peers satisfying l

$M, s \models {\sim} \langle +l \rangle^n \Box l$ iff s has more than n peers not satisfying l

$M, s \models \langle -l \rangle^n \Box {\sim} l$ iff s has at most n peers satisfying l

$M, s \models \langle +l \rangle^n \Box l$ iff s has at most n peers not satisfying l

We can abbreviate those formulas by formulas such as $[{>}\, n]x$ and $[{\leq}\, n]x$, meaning *agent has more than* n *peers satisfying* x and *agent has at most* n *peers satisfying* x, respectively, where $x \in \{p, \neg p, {\sim} p, {\sim} \neg p\}$, for $p \in At$. We can also define $[= n]x \overset{\text{def}}{=} [\leq n]x \wedge [> n-1]x$, with $n \geq 1$, meaning that *the agent has exactly* n *peers satisfying* x. For $n = 0$, define $[= 0]x \overset{\text{def}}{=} \Box {\sim} x$. Now $[= n]l \wedge [= m]{\sim} l$,

9 See [1,3,4,25] for modal logics with notions of counting.

where $l \in \{p, \neg p\}$, indicates that the agent has exactly $n + m$ peers in total. Since for any $n \in \mathbb{N}$ there are exactly $n + 1$ binary sums that equal n, we can define $[[= n]]$, meaning that an *agent has exactly n peers in total*, via a finite disjunction $([=n]p \wedge [=0]{\sim}p) \vee ([=n-1]p \wedge [=1]{\sim}p) \vee ... \vee ([=0]p \wedge [=n]{\sim}p)$.

Notice that our counting abilities are limited to $\neg$-literals (like p and $\neg p$) and their $\sim$-negations, since our base modalities $[\circ l]$ deal only with $\neg$-literals. This indicates that a consolidation taking amounts of evidence into account would have to work on the atomic level, just as our $\mathscr{C}$-consolidations, but the development of such consolidations will be left for future work.

6 Related Work

Now we briefly put our work in context with other belief formation/update theories. There are similar works, but in general our multi-agent perspective plus the qualitative and "modal" processing of evidence set our approach apart.

The term "consolidation" employed here is inspired by the homonymous belief revision operator [15,16], where an inconsistent belief base is transformed into a consistent one; likewise, our consolidations must respect the Consistency postulate. One of the most obvious differences between our approach and belief revision is that we are dealing with a multi-agent setting.

As for Bayesianism, the Bayesian update rule tells us how to update our beliefs, but not how to form them – those are the priors, which are usually allowed to be arbitrary. Our models, in principle, seem to be more in line with objective Bayesianism, which is a controversial position, but more research is needed in order to make a more rigorous comparison.

Dempster-Shafer theory of evidence [8,32] is a generalisation of probability theory where probabilities can be assigned not only to events but also to sets of events. This theory offers rules for combinations of probability assignments, which in a way can be seen as a kind of consolidation operation.

One of the main differences between our modelling and theories as Dempster-Shafer's and Bayesianism is that the latter have a clear quantitative take on evidence. Our framework employs a more limited modal language, where such quantitative statements are not even expressible (although we lay the groundwork for such possibility in Sect. 5). In our models, features such as unanimity and existence of at least one peer with some evidence play important roles, whereas in the other two theories mentioned above these notions are not straightforwardly expressible. Our paper illustrates that there are *some* sensible rationality constraints for formation of evidence-based beliefs even in a limited modal setting, but on the other hand shows the limitations of such a framework and gives the next step towards a quantitative, many-valued modal logical approach to the consolidation problem.

This modal/qualitative perspective is also one of the main differences between our models and opinion diffusion models such as [4]. Although our system is very much in the spirit of other works in opinion dynamics and aggregation and social choice theory (see e.g. [10]), our setup and treatment of evidence is unique.

This contribution does not attempt to offer a *better* formalism for multi-agent evidence-based beliefs, but to highlight how a many-valued modal logic can be used for such a task, bringing an entirely new perspective to this field.

7 Conclusions and Future Work

In this paper we took a many-valued modal logic (FVEL) and showed how it can be used to model networks of peers, where each one may have different evidence for each atomic proposition, including conflicting and incomplete evidence. We showed that in this setup, there is a question of *consolidation*: how to form beliefs given some evidence? We delineated formally a reasonable class of possible consolidations (Definition 3), using a concept similar to bisimulation. Then, we proposed postulates that have to be satisfied in order for a consolidation to be rational, and we showed that (i) they are enough to block many inadequate consolidations and (ii) they are not too strong, as they are jointly satisfiable.

Moreover, we have defined one dynamic operator with the aim of adding and removing evidence. We showed that this operator is useful to formalise some postulates inside the language, and also proposed two important new postulates formulated as axioms containing this operator, without which some unreasonable consolidations would be allowed. With these axioms, we characterised a class of consolidations satisfying most core postulates – with the exception of two which are not axiomatisable. Finally, we showed that this dynamic operator makes the language strictly more expressive, giving it the ability to "count peers", and how this lays the groundwork for quantitative consolidations that take amounts of evidence into account – but the development of those are left for future work.

A complete tableau system for FVEL is found in [29], and [27,28] give an axiomatisation for a language similar to it. Since we use a different version of FVEL, a calculus for it is still missing. Given that we already presented axioms for most postulates, an axiomatic system is preferable. It remains to be seen, however, if such axiomatisation is possible, given that some postulates are not axiomatisable and others are only "negatively" so. Considering that we have not defined a unique belief operator but only constrained the possibilities for such an operator, a complete axiomatisation for our variant of FVEL will probably require one particular consolidation to be chosen. Although we have not talked about public announcements, which in this setting are operations that *remove peers not satisfying some conditions regarding evidence, higher-order evidence or even beliefs*, we know that not all of the reduction axioms of [29] apply here.

Finally, in the consolidations presented here, the agents form beliefs based on their evidence and *their peers' evidence*. Another possibility is to make the evidence private to each agent, so that they have to resort only to their own evidence and their peers' *opinions*.

Acknowledgements. Special thanks to Barteld Kooi and Rineke Verbrugge for important suggestions. I would also like to thank the anonymous reviewers for very relevant comments. Research supported by Ammodo KNAW project "Rational Dynamics and Reasoning".

References

1. Areces, C., Hoffmann, G., Denis, A.: Modal logics with counting. In: Dawar, A., de Queiroz, R. (eds.) WoLLIC 2010. LNCS (LNAI), vol. 6188, pp. 98–109. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-13824-9_9
2. Arrow, K.J.: Social Choice and Individual Values. Wiley, New York (1951)
3. Baltag, A., Christoff, Z., Hansen, J.U., Smets, S.: Logical models of informational cascades. Stud. Logic **47**, 405–432 (2013)
4. Baltag, A., Christoff, Z., Rendsvig, R., Smets, S.: Dynamic epistemic logics of diffusion and prediction in social networks. In: Proceedings of the 12th Conference on Logic and the Foundations of Game and Decision Theory (2016)
5. Belnap, N.: A useful four-valued logic. In: Dunn J.M., Epstein G., (eds.) Modern Uses of Multiple-valued Logic, vol. 2, pp. 5–37. Springer, Dordrecht (1977). https://doi.org/10.1007/978-94-010-1161-7_2
6. Christensen, D.: Higher-order evidence. Philos. Phenomenol. Res. **81**(1), 185–215 (2010)
7. Christoff, Z., Hansen, J.U.: A logic for diffusion in social networks. J. Appl. Logic **13**(1), 48–77 (2015)
8. Dempster, A.P.: A generalization of Bayesian inference. J. Roy. Stat. Soc.: Ser. B (Methodol.) **30**(2), 205–232 (1968)
9. Elga, A.: Reflection and disagreement. Noûs **41**(3), 478–502 (2007)
10. Endriss, U., Grandi, U.: Graph aggregation. Artif. Intell. **245**, 86–114 (2017)
11. Fagin, R., Halpern, J.: Belief, awareness, and limited reasoning. Artif. Intell. **34**(1), 39–76 (1987)
12. Foley, R.: Self-trust and the authority of others. In: Foley, R. (ed.) Intellectual Trust in Oneself and Others, pp. 83–130. Cambridge University Press, Cambridge (2001)
13. Fricker, E.: Testimony and epistemic autonomy. In: Lackey, J., Sosa, E. (eds.) The Epistemology of Testimony, pp. 225–250. Oxford University Press, Oxford (2006)
14. Gibbard, A.: Manipulation of voting schemes: a general result. Econometrica **41**(4), 587–601 (1973)
15. Hansson, S.O.: Belief base dynamics. Ph.D. thesis, Uppsala University (1991)
16. Hansson, S.O.: Semi-revision. J. Appl. Non-Class. Logics **7**(1–2), 151–175 (1997)
17. Hardwig, J.: Epistemic dependence. J. Philos. **82**(7), 335–349 (1985)
18. Heyting, A.: Intuitionism: an Introduction, vol. 41. Elsevier (1966)
19. Hintikka, J.: Impossible possible worlds vindicated. In: Saarinen, E. (ed.) Game-Theoretical Semantics, vol. 5, pp. 367–379. Springer, Dordrecht (1979). https://doi.org/10.1007/978-1-4020-4108-2_13
20. Hintikka, J.: Knowledge and Belief: An Introduction to the Logic of the Two Notions, Contemporary Philosophy, vol. 4. Cornell University Press, Ithaca (1962)
21. Kelly, T.: Peer disagreement and higher order evidence. In: Goldman, A.I., Whitcomb, D. (eds.) Social Epistemology: Essential Readings, pp. 183–217. Oxford University Press (2010)
22. Lasonen-Aarnio, M.: Higher-order evidence and the limits of defeat. Philos. Phenomenol. Res. **88**(2), 314–345 (2014)
23. Lehrer, K.: Social information. The Monist **60**(4), 473–487 (1977)
24. Martini, C., Sprenger, J., Colyvan, M.: Resolving disagreement through mutual respect. Erkenntnis **78**(4), 881–898 (2013)
25. Pacuit, E., Salame, S.: Majority logic. In: Dubois, D., Welty, C., Williams, M.A. (eds.) Principles of Knowledge Representation and Reasoning, KR 2004, vol. 4, pp. 598–605. AAAI Press (2004)

26. Rantala, V.: URN models: a new kind of non-standard model for first-order logic. J. Symbolic Logic **4**, 455–474 (1975)
27. Rivieccio, U.: Algebraic semantics for bilattice public announcement logic. In: Studia Logica Proceedings Trends in Logic XIII. Springer, Heidelberg (2014)
28. Rivieccio, U.: Bilattice public announcement logic. In: Goré, R., Kooi, B., Kurucz, A. (eds.) AiML, vol. 10, pp. 459–477. College Publications (2014)
29. Santos, Y.D.: A dynamic informational-epistemic logic. In: Madeira, A., Benevides, M. (eds.) DALI 2017. LNCS, vol. 10669, pp. 64–81. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-73579-5_5
30. Santos, Y.D.: Consolidation of belief in two logics of evidence. In: Blackburn, P., Lorini, E., Guo, M. (eds.) LORI 2019. LNCS, vol. 11813, pp. 57–70. Springer, Heidelberg (2019). https://doi.org/10.1007/978-3-662-60292-8_5
31. Satterthwaite, M.A.: Strategy-proofness and arrow's conditions: existence and correspondence theorems for voting procedures and social welfare functions. J. Econ. Theory **10**(2), 187–217 (1975)
32. Shafer, G.: A Mathematical Theory of Evidence, vol. 42. Princeton University Press, Princeton (1976)
33. Troelstra, A.S., Van Dalen, D.: Constructivism in Mathematics, vol. 1. North Holland, Amsterdam (1988)