

Belief Integration and Source Reliability Assessment

Paolo Liberatore

*DIAG – Sapienza University of Rome
Via Ariosto 25, 00185, Rome, Italy*

PAOLO@LIBERATORE.ORG

Abstract

Merging beliefs requires the plausibility of the sources of the information to be merged. They are typically assumed equally reliable when nothing suggests otherwise. A recent line of research has spun from the idea of deriving this information from the revision process itself. In particular, the history of previous revisions and previous merging examples provide information for performing subsequent merging operations.

Yet, no examples or previous revisions may be available. In spite of the apparent lack of information, something can still be inferred by a try-and-check approach: a relative reliability ordering is assumed, the sources are integrated according to it and the result is compared with the original information. The final check may contradict the original ordering, like when the result of merging implies the negation of a formula coming from a source initially assumed reliable, or it implies a formula coming from a source assumed unreliable. In such cases, the reliability ordering assumed in the first place can be excluded from consideration.

Such a scenario is proved real under the classifications of source reliability and definitions of belief integration considered in this article: sources divided in two, three or multiple reliability classes; integration is mostly by maximal consistent subsets but also weighted distance is considered. Other results mainly concern the integration by maximal consistent subsets and partitions of two and three reliability classes.

1. Introduction

Between November 2005 and September 2006 Wikipedia had an article about Porchesia, a 300,000-inhabitant island in the Mediterranean sea (Wikipedia, 2006). No such place was ever mentioned in the article on the Mediterranean sea or those of Europe, Asia or Africa. It took ten months for the article to be found out a hoax and removed.

Wikipedia accepts information from any user who wants to provide some, even anonymously; no prior barrier exists, not even for article creation. This was the reason of its success: a previous project of creating an on-line encyclopedia from reputed experts failed (Wiley & Gurrell, 2009). At the same time, it opens the door to wrong and malicious information. This is not limited to Wikipedia. A number of other successful sites work similarly: computer programmers rely on Stack Overflow (Barua, Thomas, & Hassan, 2012), information on many topics can be found on Stack Exchange (Posnett, Warburg, Devanbu, & Filkov, 2012) and Yahoo Answers (Adamic, Zhang, Bakshy, & Ackerman, 2008). Everyone can provide information; checking it is done later. Wikipedia has article removal and correction to fix incorrect additions, other Internet fora has measures of trust.

The Porchesia article is an extreme example, but shows the difficulty of manually checking a large amount of information coming from unknown sources, even by a large and dedicated community (Viégas, Wattenberg, Kriss, & van Ham, 2007). A large island should

have been at least mentioned in related articles, such as that about the Mediterranean sea, but it was not. But these articles also come from users of unknown reliability. The conflict is between sources of unknown reliability.

Unknown does not mean the same for all, as commonly assumed in the theory of belief merging (Konieczny & Pérez, 2011; Liberatore & Schaerf, 1998; Everaere, Konieczny, & Marquis, 2015; Cevolani, 2014; Herzig, Pozos-Parra, & Schwarzentruher, 2014; D’Alfonso, 2016). The user who created the hoax article is not to be trusted much when creating other new articles on Wikipedia.

Assuming a prior assessment of the sources would be unrealistic, and not only for Web sites such as Wikipedia and Stack Overflow, which rely on new users providing useful information. Even in more controlled scenarios, such as database integration, assigning priority to sources is only one of several value conflict strategies (Naumann, Bilke, Bleiholder, & Weis, 2006), the others including “ask the user”, “take the most used value”, “if the value is numeric, use the average”. Some works in belief revision try to derive such a preference from the previous history of revisions (Booth & Nittka, 2008; Liberatore, 2015) or from examples (Liberatore, 2016). Neither is assumed known in this article.

The scenario considered in this article comprises some sources of information, each providing information expressed as some propositional formulae. In the simplest case each source gives exactly one formula. More generally, several pieces of information may come from the same source: a sensor produces several readings, a database can be queried several times, etc.




Figure 1: Multiple sources providing formulae to be merged

An example of such settings is in Figure 1: F_1 , F_3 and F_5 come from the first source, F_4 and F_7 from the second, F_2 , F_6 and F_8 from the third. The sources may be differently reliable; for example, the first is more reliable than the third. The result of merging depends on the reliability of the sources; for example, if $F_1 = x$, $F_2 = \neg x$ and none of the other formulae mention x , the result of merging should contain x since the source supporting x is more reliable than the source denying it. In general, the merger gives preference to formulae coming from reliable sources. Reliability is attached to the sources, and is therefore the same for all formulae provided by the same source. The case of different reliability in the same source is discussed separately.

The relative reliability of the sources may not be available. A large amount of work done assuming equal reliability (Konieczny & Pérez, 2011; Liberatore & Schaerf, 1998; Everaere

et al., 2015; Cevolani, 2014; Herzig et al., 2014; D’Alfonso, 2016), and how to obtain reliability information has been largely overlooked. As observed by Konieczny (2004):

If one has some notion of relative reliability between sources, it is enough and sensible to force the less reliable ones to give up first [...] But often we do not have such information [...]

A recent line of research tries to address this problem by deriving such metainformation from the merging process itself: Booth and Nittka (2008) and Liberatore (2015) used the history of previous revisions to perform the following revisions; Liberatore (2016) exploited merging examples; Booth (2006) and Konieczny (2004) weaken the sources that are furthest from the others.

In lack of an explicit measure of source reliability, a previous history or a set of merging examples, it may seem that the only sensible policy is to assume that sources are equally reliable. Yet, something more can be said even in this case from the very definition of reliability.

Reliability tells how much each formula affects the result of merging. If the most reliable source says x and the least says $\neg x$, the result should imply x or at least be consistent with it. Otherwise, merging did not really reflect the assumed reliability of the sources. In the other way around, the given reliability is not coherent with the result of merging it produces.

This observation can be used as a method for testing a candidate measure of reliability: it is used in a merging attempt, and if the result is not coherent with the candidate reliability ordering this candidate is discarded. The typical application is with a specific source of merging that is totally reliable: every piece of information it provides is sure, albeit possibly incomplete; another source providing information that contradicts it cannot be considered reliable. Still better, the result of merging can be compared with the other sources to assess their reliability. This allows for excluding some reliability orderings even in lack of a totally reliable source.

Such a scenario is depicted in Figure 2. The process comprises three steps: first, a reliability ordering of the sources is considered, in this case that sources A and B are more reliable than C (this ordering is formally represented by $(AB|C)$, at the left of the figure); second, the formulae provided by the three sources $A = \{x\}$, $B = \{\neg x, \neg y\}$ and $C = \{x, \neg y\}$ are combined based on this ordering (box “merger” in the figure, in this case by maximal consistent subsets); one possible outcome is $\{x, \neg y\}$; this set is compared with the sources in the third step (box “evaluator”); since A and C are fully consistent with this set while B is not, A and C can be considered more reliable than B . This contradicts the assumption that B is more reliable than C . The initial reliability ordering can be excluded.

The initial reliability ordering can be excluded because using it for merging produces a result that assesses the reliability of the sources differently from itself. Coherence between what was assumed and what is concluded is a minimal requirement for an accepting an assumption.

Example 1 *The mayor of a mountain town has to quickly decide how to secure some houses in danger of a landslide. He does not want to be blamed for the likely disaster, so he collects opinions from three experts: Akai, Bialetti and Connors. To his dismay, their reports are partly conflicting as to the cause and possible effects of the landslide and do not*

therefore agree on the action to take. He has to take a decision in a short time to avoid looking undecided, but the responsibility of its effects will be his.

He calls up a meeting with his supporters in the city council, but they do not understand anything about disaster prevention in general and landslides in particular. Instead, they offer political advice: publicly praise the experts for their help, and state that the decision is based on their opinions. This way, he can later blame them in case of failure. Since the opinions of the experts diverge, whichever decision is taken will conflict with some of them; yet, he can still make the excuse that some experts are less reliable than the others.

As for which experts are “less reliable”, they are at loss. Worse, none wants to be later accused of suggesting a choice that resulted in a political disaster. The mayor tosses a coin: Akai and Bialetti are his reputed experts, Connors not that much. Some other tosses, and a decision based on merging Akai’s and Bialetti’s opinions is taken: build a retaining wall.

A week later, the landslide kills two dozen workers and buries twelve houses with their inhabitants inside.

Fortunately, the mayor and his supporters have their plan to divert responsibility from them to the two reputed experts. At the press conference, the mayor mourns the dead and blames Akai and Bialetti for their ill advice. In spite of their international reputation proved by awards and invitations to conferences, they were wrong.

When a journalist asks about Connors, the mayor is ready to answer with his excuse, suggested by one of his assistants five minutes before the press conference: Connors gave many contributions to geology but now he is old, some of his knowledge of the field might be outdated now.

So why was the wall built, if this choice was consistent with Connors’ opinion? And why, if some parts of Bialetti’s report suggested the wall was useless? The questions echo from the press conference to the newspapers and the local TVs in the following days. Excerpts from Bialetti’s and Connors’ reports are repeatedly shown; they prove that the decision purportedly based on the first and distrusting the second was in fact inconsistent with the first and consistent with the second.

The mayor ends up looking so dumb to not even be able to follow the suggestions of the experts he himself chose, and is forced to resign.

One of the next sections shows that a scenario like this is indeed possible: some ordering of the reliability of the experts (the sources) leads to a choice (building the wall) that contradicts the ordering itself. The question therefore arises: how could the mayor have avoided this outcome and instead remain firmly in office in spite of his incompetence and cowardice?

This example shows how the overall process works: in lack of knowledge of reliability, every order of reliability could potentially be assumed (for example, Akai and Bialetti are more reliable than Connors); however, the result of merging (which supports building the retaining wall) may then conflict with the original assumption (building a retaining wall is consistent with Connors’ report but not with Bialetti’s). Consistency between what is assumed and what is concluded is a minimal requirement for the validity of the assumption.

The coherence between what was assumed and what is concluded is a minimal requirement for an assumption to be acceptable. For this reason, fixpoint definitions are widespread in nonmonotonic reasoning: for example, an extension of a default theory (Reiter, 1980)

is a set of formulae that can be derived from the premises using rules that depend on the extension itself. This article applies this common approach to the reliability of the information sources instead of single atoms as in answer set programming (Bonatti, Calimeri, Leone, & Ricca, 2010) or formulae like in autoepistemic logic (Gabbay & Schlechta, 2016).

The example fits into the formal models of sources and formulae: each report is a source, each of its statements one of the propositional formulae in the source. Akai's report tells that the subsoil is stable at the toe of the landslide; Bialetti instead wrote that it is not stable, but adds that the landslide mass is not large; Connors says that the subsoil is stable and confirms that the landslide mass is not large. In formulae, this is $A = \{x\}$, $B = \{\neg x, \neg y\}$ and $C = \{x, \neg y\}$, where x is the stability of the subsoil at the toe of the landslide and y the largeness of the landslide. One way to consistently integrate $(AB|C)$ is to accept x from A and $\neg y$ from B : the subsoil is stable and the landslide mass not large, meaning that a gravity retaining wall can be built and will be effective at stopping the landslide (Seattle DPD, 2000). The final decision to build the wall is fully consistent with C (in spite of its assumed unreliability) but not completely consistent with B (in spite of its assumed reliability): the resulting reliability ordering is $(AC|B)$.

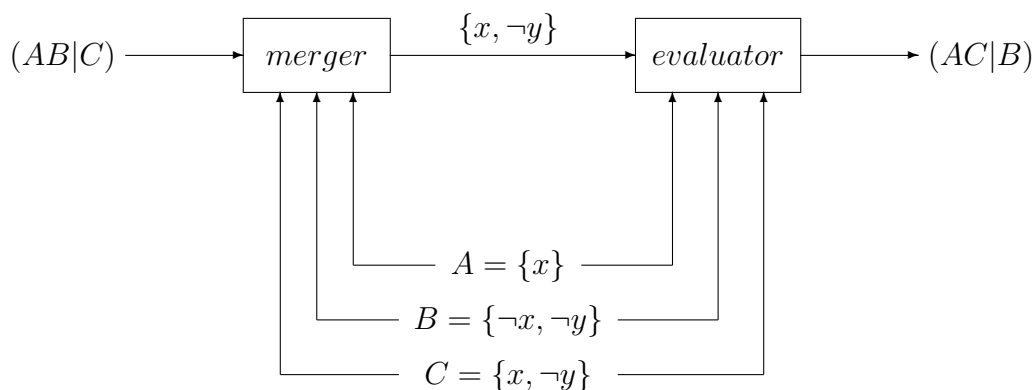


Figure 2: A reliability ordering to exclude

The example shows that the result of the process is not just a formula (the result of merging). It is also a method to establish the coherence of a reliability ordering, and therefore to discriminate among all possible reliability orderings. More generally, it is a mechanism for offering a rational justification of the merging process. As the example shows, sometimes the ability to back up a choice is key to success more than the validity of the choice.

The framework can be instantiated to different reliability orderings and integration methods. For example, the sources can be partitioned into two or three reliability classes; the results obtained in the two cases are different. While most of the results in this article are about integrating by maximal consistent subsets (Rescher, 1964; Brewka, 1989; Nebel, 1992; Baral, Kraus, Minker, & Subrahmanian, 1992; Nebel, 1998; Rott, 1993; Delgrande, Dubois, & Lang, 2006; Benferhat & Yahi, 2009), the approach can be carried to other methods, for example merging by weighted distance (Revesz, 1997; Konieczny & Pérez,

2011; Konieczny, Lang, & Marquis, 2002, 2004), some preliminary results being presented in Section 6.3. Here is a roadmap of the next sections:

- general definitions (reliability ordering, maxcons): Section 2;
- the framework is applied to the simple case in which sources are assessed as either reliable or not reliable, and each provides a single formula: Section 3;
- the constraint that sources provide a single formula each is lifted: Section 4;
- the sources are assessed as either reliable, partly reliable and unreliable: Section 5;
- the sources are partitioned according to the percentage of correct formulae they provide, or their reliability is represented as an arbitrary order, or weighted merge is used: Section 6;
- discussion of results, comparison with related work, future directions, recap of technical results and conclusions: Sections 7, 8, 9 and 10.

2. Sources and Maxcons

In this article, a source is identified with the information it provides: a source is a set of formulae. The example in Figure 1 has three sources. The first provides formulae F_1 , F_3 and F_5 , and is therefore formalized by the set $S_1 = \{F_1, F_3, F_5\}$; the second source is $S_2 = \{F_4, F_7\}$, the third $S_3 = \{F_2, F_6, F_8\}$.

Every source is assumed to provide information with the same degree of reliability; a “physical” source giving formulae at two different levels of reliability is captured as two separate sources. In the example of the mayor, a press expert that suggests what to say on television, how to react to criticism from the newspapers and which kind of wall will save the houses from the landslide would be regarded as two sources: the professional gives the first two statements, the amateur the latter. The issue of separating the information coming from a source has been recently investigated by Booth and Hunter (2018).

The relative reliability of the sources is a total preorder: every two elements compare in a way or another, but can also be equal. Such an ordering is the same as an ordered partition of the set of sources (Stanley, 1997): for example, the situation in which the sources S_1 and S_2 are more reliable than S_3 is represented by $(S_1S_2|S_3)$. The first class of the partition contains the two sets S_1 and S_2 , the second S_3 . Writing their formulae explicitly, the partition is:

$$(\{F_1, F_3, F_5\}, \{F_4, F_7\}|\{F_2, F_6, F_8\})$$

While $(S_1S_2|S_3)$ contains three sources, the similar partition $(S_1 \cup S_2|S_3)$ only contains two. Yet, the formulae and their reliability are the same: once sources are framed into an ordered partition, merging can be done by flattening each class into a single set of formulae. In the example:

$$\begin{array}{ll} \text{partition:} & (\{F_1, F_3, F_5\}, \{F_4, F_7\}|\{F_2, F_6, F_8\}) \\ \text{flattening:} & (F_1F_3F_4F_5F_7|F_2F_6F_8) \end{array}$$

While the concept of flattening should be obvious from the example, its formal definition is given for reference.

Definition 1 *The flattening of a partition $(SS_1|SS_2|\dots|SS_m)$, where each SS_i is a set of sets of formulae, is the partition of formulae in which each class contains all formulae in the sets of the corresponding class: $(S_1|S_2|\dots|S_m)$ where $S_i = \cup SS_i$.*

The flattening of a partition is a sequence of sets of formulae similar to those used in prioritized belief revision and contraction (Brewka, 1989; Rott, 1993) and in iterated belief revision (Rott, 2009), which is in fact the same as belief merging from sources of differing reliability (Delgrande et al., 2006).

On such sequences, maxcons are a way to derive consistent integrations (others are considered in following sections). The most common definition is for a single-class sequence (P) or, equivalently, a set of formulae P :

Definition 2 *A maxcon of a set of formulae P is a consistent subset of P such that no other consistent subset of P strictly contains it.*

The word “maxcon” is an abbreviation for *maximal consistent subset*. If P is consistent, its only maxcon is P itself. Otherwise, P can be made consistent by removing some of its formulae. This can be done in various ways, leading to different maxcons. For example, $\{x, \neg y, y\}$ can be made consistent by removing either $\neg y$ or y , leading to the two maxcons $\{x, y\}$ and $\{x, \neg y\}$. Maximality ensures that not too much is removed: for example, $\{x\}$ is also a consistent subset, but deleting both y and $\neg y$ for restoring consistency is an overkill.

The maxcons of a single-element sequence (P) are the maxcons of the set P . A multiple-element sequence $(P_1|P_2|\dots|P_m)$ represents an ordering of preference over formulae in which P_1 ranks first, P_2 second, and so on. If $P_1 \cup P_2 \cup \dots \cup P_m$ is consistent this ordering does not matter, otherwise it tells which formulae to remove to restore consistency. For example, $(\neg y|x, y)$ means that $\neg y$ is to be preferred over both x and y ; therefore, the best way to restore consistency is to remove y , obtaining $\{\neg y, x\}$; the second formula x is retained, even if it was in the second class of the partition, because it was not involved in the contradiction; in other words, its removal would not help restore satisfiability. The principle is: select as many consistent formulae as possible, starting from the first class (Rescher, 1964; Brewka, 1989; Nebel, 1992, 1998; Rott, 1993; Delgrande et al., 2006).

Definition 3 *A maxcon of a sequence $(P_1|\dots|P_m)$ is a subset $M \subseteq P_1 \cup \dots \cup P_m$ such that $M \cap (P_1 \cup \dots \cup P_i)$ is a maxcon of $P_1 \cup \dots \cup P_i$ for all $1 \leq i \leq m$.*

By virtue of the universal quantifier over i , $M \cap P_1$ needs to be a maxcon of P_1 . Only when this condition is met the check with P_2 matters. This ensures that formulae of the first class are removed only if this is really necessary to remove contradiction, even disregarding all other formulae. This way, maximality within P_1 takes precedence over maximality within P_2 , which then take precedence over maximality within P_3 , etc.

Definition 4 *A maxcon of a partition of sources is a maxcon of its flattening.*

The reliability of a formula is the same as the reliability of the source it comes from. It is this comparison of the formulae that the merging process uses to integrate them. From this point of view, a partition like $(\{F_1, F_3, F_5\}, \{F_4, F_7\} | \{F_2, F_6, F_8\})$ is the same as its flattening $(F_1 F_2 F_3 F_5 F_7 | F_4 F_6 F_8)$ since the reliability of the formulae is the same: F_1 is as reliable as F_3 , both are more so than F_2 , etc.

Collapsing $(P_1 | \dots | P_m)$ into $P_1 \cup \dots \cup P_m$ makes all formulae compare the same. The maxcons of this set are called the plain maxcons of the partition.

Definition 5 *The plain maxcons of a partition of sources are the maxcons of the sets of all formulae in the sources.*

Every maxcon is also a plain maxcon, but not the other way around. Indeed, plain maxcons are obtained by first deleting the ordering over formulae; therefore, some plain maxcons may remove formulae that occur early in the partition and keep others that occur late.

Some results carry over from maxcons of sets to maxcons of partitions. The first is an almost direct consequence of being a maximal subset: it contains every formula consistent with it.

Lemma 1 *If M is a maxcon of $(P_1 | \dots | P_m)$ and M is consistent with $F \in P_1 \cup \dots \cup P_m$, then M contains F .*

Proof. By definition, $M \cap (P_1 \cup \dots \cup P_i)$ is a maxcon of $P_1 \cup \dots \cup P_i$ for every $i \in \{1, \dots, m\}$; for $i = m$ this condition is that M is a maxcon of $P_1 \cup \dots \cup P_m$. If $F \notin M$ but $M \cup \{F\}$ is consistent, then M is not a maximal consistent subset of $P_1 \cup \dots \cup P_m$. \square

A maxcon of a partition is also a maxcon of the set comprising all formulae in the partition. The converse does not hold in general. For example, the partition $(x | \neg x)$ has a single maxcon $\{x\}$, but the set $\{x, \neg x\}$ has the two maxcons $\{x\}$ and $\{\neg x\}$. The second is not a maxcon of the partition because it prefers the lower-ranked formula $\neg x$ over the higher-ranked x .

If a set is consistent, its only maxcon is itself. A similar result holds for partitions in different ways: for P_1 only, for $P_1 \cup \dots \cup P_m$, or for every union in between.

Lemma 2 *If $P_1 \cup \dots \cup P_i$ is consistent, every maxcon of $(P_1 | \dots | P_m)$ is a superset of $P_1 \cup \dots \cup P_i$.*

Proof. By definition of maxcons of a partition, $M \cap (P_1 \cup \dots \cup P_i)$ is a maxcon of $P_1 \cup \dots \cup P_i$. Since this union is consistent, its only maxcon is itself: $M \cap (P_1 \cup \dots \cup P_i) = P_1 \cup \dots \cup P_i$. This implies that M contains all of $P_1 \cup \dots \cup P_i$ \square

The converse holds for partitions made of two classes.

Lemma 3 *If a maxcon of $P_1 \cup P_2$ is a superset of P_1 , it is a maxcon of $(P_1 | P_2)$.*

Proof. Let M be a maxcon of $P_1 \cup P_2$ such that $P_1 \subseteq M$. By definition, M is a maxcon of $(P_1|P_2)$ if it is a maxcon of $P_1 \cup P_2$ and $M \cap P_1$ is a maxcon of P_1 . The first condition holds by assumption. The second holds as well: M being a maxcon, it is consistent; therefore, its subset $M \cap P_1$ is consistent; it is maximally consistent within P_1 because $M \cap P_1 = P_1$. \square

Given a partition $(P_1 | \dots | P_m)$, its plain maxcons are obtained by flattening the partition into a set $P_1 \cup \dots \cup P_m$. A maxcon is also a plain maxcon; the converse may hold or not. It holds for two-class partitions under a simple condition.

Lemma 4 *A plain maxcon M of a partition $(P_1|P_2)$ is also a maxcon of the same partition if no other plain maxcon M' is such that $M \cap P_1 \subset M' \cap P_1$.*

Proof. For two-classes partitions, the definition of maxcons specializes to: $M \cap P_1$ is a maxcon of P_1 and M is a maxcon of $P_1 \cup P_2$. The second part is the same as the definition of a plain maxcon. The first part is ensured by the statement of the lemma. \square

The case of partitions of two classes is the simplest considered in this article: sources are assessed as fully reliable or not. For example, if the source providing S_1 and S_3 are reliable while S_2 is unreliable, the partition is $(S_1 S_3 | S_2)$. With the abuse of notation announced above, this is the same as $(S_1 \cup S_3 | S_2)$.

One of the following theorems relies on three formulae A and B such that their maxcons are $\{A, B\}$ and $\{B, C\}$. Once such formulae are provided, the rest of the proof is relatively simple. More importantly, the specific formulae used for A , B , and C no longer matter, but only that their maxcons are $\{A, B\}$ and $\{B, C\}$. The proof comprises two separate and almost disconnected steps: a. first, build three formulae with the required maxcons (for example, $A = x \wedge \neg y$, $B = x \vee y$, and $C = \neg x \wedge y$); and b. use the fact that their maxcons are $\{A, B\}$ and $\{B, C\}$ to prove the claim.

Such a scheme emerges in several other proofs. The first part is common to all of them: certain formulae F_1, \dots, F_m have a given set of maxcons M_1, \dots, M_k ; the second part of the proof infers the claim from this fact. In some cases, the formulae and the maxcons are many, and explicitly writing the formulae and proving that their maxcons are as required is long and tedious. The following lemma comes to help: it shows that this part of these proofs can be taken for granted.

It tells that one can simply say “some formulae F_1, \dots, F_m have maxcons M_1, \dots, M_k ” without explicitly writing these formulae and determine their maxcons. The only requirement is that none of these maxcons is contained in another. For example, the proof using three formulae A , B , and C with maxcons $\{A, B\}$ and $\{B, C\}$ needs not to explicit A , B and C : since the two required maxcons $\{A, B\}$ and $\{B, C\}$ are not contained one in the other, such formulae exist.

Lemma 5 *Given some sets of symbols, none contained in another, there exists a formula for each symbol so that the maxcons of these formulae correspond to the given sets of symbols.*

Proof. The proof was previously published (Liberatore, 2016), but is restated here for completeness.

For n sets, $\lceil \log n \rceil$ propositional variables are required. Each set of symbols is associated with a unique propositional interpretation; this is possible because by construction there are at least n propositional interpretations over these variables.

For each such interpretation, there exists a formula that is satisfied only by it. For example, if the interpretation makes x and y false and z true, the formula is $\neg x \wedge \neg y \wedge z$. Since each set of symbols is associated with a propositional interpretation, is also associated with the corresponding formula.

If symbol L is in the sets S_1, S_2, \dots , and these sets corresponds to formulae F_1, F_2, \dots , the formula of L is their disjunction $F_1 \vee F_2 \vee \dots$. As a result, the formula corresponding to the symbol L is satisfied exactly by the interpretations of the sets S_1, S_2, \dots .

By construction, if a set of symbols is associated with the interpretation I , the formulae corresponding to the symbols in the set are satisfied by I . This proves that each set of symbols corresponds to a consistent set of formulae. This set is also maximally consistent because: a. no other formula is satisfied by that interpretation; and b. if all formulae of the set plus some others are satisfied by another interpretation, then the set corresponding to that interpretation strictly contains the considered one, contradicting the assumption that none of the sets strictly contains another.

To conclude the proof, the formulae do not have other maxcons. This is because the formulae are only satisfied by some of the interpretations corresponding to the sets of symbols, and each of them is the only model of a maxcon. \square

3. Two-Class Partitions of Singleton Sources

In the example of the mayor and the landslide, three experts provide conflicting information, and their reliability is unknown. Therefore, every ordering of reliability of them is potentially valid, even one chosen by tossing a coin. The problem with the particular one used by the mayor was not its randomness, since he could easily purport a reason for backing it up. The problem was that the decision based on this order was later found to be inconsistent with the order itself.

The question is therefore: is a situation like this really possible? If so, can one just exclude such self-conflicting orderings and choose one among the other ones?

To simplify the exposition, this section analyzes the case where each source provides exactly one piece of information, encoded a single propositional formula. The reliability ordering is a simple yes/no classification: sources are either fully reliable or not; in the example of the landslide, the mayor decided that Akai and Bialetti are reliable, Connors cannot be fully trusted. Both assumptions are released in the following sections; yet, this very simple case already allows for some significant results.

Given a classification of the sources, each maxcon of the ordered partition is a way for merging them while retaining as much of them as consistency allows. If the sources $S_1 = \{x, y \vee z\}$, $S_2 = \{\neg y \wedge \neg z\}$ and $S_3 = \{x \wedge y\}$ are partitioned as $(S_1 S_2 | S_3)$, their maxcons are $\{x, y \vee z, x \wedge y\}$ and $\{x, \neg y \wedge \neg z\}$. Each maxcon is a way of consistently selecting the formulae from the most reliable ones to the least. In this case, the most reliable formulae are x , $y \vee z$ and $\neg y \wedge \neg z$; the latter two contradict each other, so they are alternatives; the first alternative can be consistently added $\neg x \wedge \neg y$. These two alternatives are the two maxcons.

This process hinges on the relative reliability of the formulae. What if this information is not given? What if the reliability of the formulae is not known? Formula x may be more reliable than $\neg x$. Or less. Or equally. The third is only one possibility. Yet, equal reliability is often assumed in lack of information telling otherwise. This “equality by ignorance” principle is often assumed not only in belief revision but also in nonmonotonic reasoning and in probability theory: unprioritized conflicting defaults are taken as alternatives (Katz & Golbeck, 2006), unknown probabilities are assumed equal according to the much-criticized principle of indifference (Shackel, 2007).

In terms of ordered partitions, equal reliability is encoded by the single-class ordered partition $(S_1 S_2 S_3)$. But nothing rules out any of the other partitions, for example $(S_1 S_2 | S_3)$. Still better, if reliability is unknown then S_1 and S_2 may be reliable the same and more so than S_3 , as far as it is known. This is a sensible possibility. Then merging enters the picture: a maxcon like $\{x, y \vee z, x \wedge y\}$ is a way to combine the formulae according to the reliability encoded by $(S_1 S_2 | S_3)$. It derives from the sources and their assumed reliability. Checking this maxcon with the formulae reveals a problem: S_3 is consistent with it while S_2 is not. This is a problem because S_2 was assumed more reliable than S_3 , but the conclusion supports the contrary. Different partitions such as $(S_1 S_2 | S_3)$ do not fall into this trap: its maxcon is coherent with the reliability it encodes. Partitions like this are consistent with themselves.

The following sections use the same principle, which in this simple case can be summarized as:

1. assume an arbitrary partition of the sources (reliability ordering);
2. determine a possible consistent combination of the formulae they provide (a maxcon);
3. use it to classify the sources: the ones consistent with it are reliable, the others are unreliable;
4. check if this classification is the same as that initially assumed.

This section assumes that sources are singletons (each source provides exactly one formula) and that partitions have two reliability classes. The following sections relax both constraints, but the general principle stays the same: assume some relative reliability of the sources (1), combine their formulae (2), use this combination for assessing the reliability of the sources (3); if this is not the same as that originally assumed, discard it (4).

Assuming a partition of the sources, each maxcon is a way to combine them according to this ordering while avoiding inconsistency. What is missing is the third step of the process: from a consistent combination of formulae, assess the reliability of the sources. This is done by checking the consistency of the sources with each maxcon M .

Definition 6 *A set of formulae M induces on the set of sources S the partition $(R|S \setminus R)$, where:*

$$R = \{S_i \mid M \not\models \neg A_i \text{ for all } A \in S_i\}$$

This definition closes the circle: from a partition one can determine its maxcons, and each maxcon induces a partition of the sources. Technicalities apart, assuming a reliability ordering of the sources allows combining them; this combination allows assessing the reliability of the sources.

Definition 7 *A partition is stable if it has a maxcon whose induced partition is the partition itself.*

Stability could also be defined by requiring all maxcons to induce the partition itself. This stronger condition is used in the next sections. For this introductory part of the article, the weaker definition is used: if there is a possible way to perform the merging (a maxcon) such that the result is consistent with a source, the source is considered reliable. The results obtained from this definition are:

- stable partitions always exist;
- not all partitions are stable;
- the maxcons of the stable partitions are the plain maxcons.

The first two facts show that the definition of stability is not trivial: some reliability orderings lead to combinations of formulae that contradict them and some others do not; the constraint of stability gives some useful information in removing some but not all of them from consideration. The third fact may look disappointing: even with the condition of stability all plain maxcons are still obtained. In the next sections the assumption of singleton sources and binary partitions will be lifted, and this result no longer holds. For now, it can be seen as a confirmation that in the basic case the mechanism is sensible: lacking other formulae coming from the same source, the reliability of a formula depends only on itself; therefore, nothing allows assuming a formula more reliable than another.

Theorem 1 *There exist three formulae A , B and C such that the partition $(AB|C)$ is stable while $(A|BC)$ is not.*

Proof. Three consistent formulae A , B and C are chosen so that their plain maxcons are $\{A, B\}$ and $\{B, C\}$. Lemma 5 proves that this is possible because these sets are not contained in each other; for example, $A = (x \vee z) \wedge y$, $B = x \wedge \neg z$ and $C = (x \vee z) \wedge \neg y$.

Formulae with such maxcons can be graphically depicted in the space of models as in Figure 3. Each formula is a box representing its set of models and intersections indicate mutual consistency.

The partition $(AB|C)$ is stable. Since $\{A, B\}$ is consistent but conflicts with C , it is the only maxcon of the partition. The reliability of the sources can be assessed from it. Since A and B are consistent with $\{A, B\}$, their sources $\{A\}$ and $\{B\}$ are in the first class of the partition induced by the maxcon. Instead, C is inconsistent with the maxcon; therefore, its source $\{C\}$ is in the second class of the induced partition. The resulting partition is therefore $(AB|C)$, the same assumed in the first place.

The partition $(A|BC)$ is not stable. The plain maxcon $\{B, C\}$ is excluded because it does not contain the most reliable formula A . The only plain maxcon that is also a maxcon




Figure 3: The formulae in the proof of Theorem 1, depicted in the space of models. Since A and B are consistent their sets intersect. The same for B and C .

of the partition is $\{A, B\}$. This set is consistent with A and B but not with C . Therefore, the induced partition is $(AB|C)$, while that initially assumed was $(A|BC)$. \square

The maxcons of A , B and C are $\{A, B\}$ and $\{B, C\}$. The first is generated by the stable partition $(AB|C)$; by symmetry, the second is generated by $(BC|A)$. In this particular case, both maxcons are generated by some stable partition. More generally, this happens whenever all sources contain a single formula and partitions comprise two classes.

Lemma 6 *If S is a set of singleton sources, $R \subseteq S$ and $(R|S \setminus R)$ is a stable partition, then $\cup R$ is its only maxcon.*

Proof. By definition of stability, $(R|S \setminus R)$ has a maxcon M such that R are the sources consistent with M . Since every formula F in R is consistent with M , Lemma 1 implies $F \in M$. In other words, $\cup R \subseteq M$.

Since R is the set of sources that are consistent with M , all other sources in S are inconsistent with M , and are therefore not in M . This implies that $M = \cup R$.

Let M' be an arbitrary maxcon of $(R|S \setminus R)$. Since $\cup R = M$ and M is consistent by definition of maxcon, so is $\cup R$. By Lemma 2, every maxcon of $(R|S \setminus R)$ is a superset of $\cup R$, that is, $\cup R \subseteq M'$. But no source in $S \setminus R$ is consistent with $M = \cup R$, implying that none is consistent with M' . This proves that $M' = M$. \square

This lemma is used to prove that the stable partitions generate exactly the plain maxcons of the set of sources if sources are singletons and partitions comprise two classes.

Theorem 2 *For every $R \subseteq S$ where S is a set of singleton sources, $\cup R$ is a maxcon of $\cup S$ if and only if $(R|S \setminus R)$ is stable.*

Proof. By Definition 4, a maxcon of a partition of sources is a maxcon of its flattening. Therefore, $\cup R$ is a maxcon of $(R|S \setminus R)$ if and only if it is a maxcon of $(\cup R | \cup (S \setminus R))$. In turn, Definition 3 formulates this as (a) $\cup R$ is a maxcon of $(\cup R) \cup (\cup (S \setminus R)) = \cup S$ and (b) $(\cup R) \cap (\cup R) = \cup R$ is a maxcon of $\cup R$.

Let $(R|S \setminus R)$ be a stable partition. By Lemma 6, $\cup R$ is its only maxcon. Therefore, Condition (a) is true, and $\cup R$ is a maxcon of $\cup S$.

Let $\cup R$ be a maxcon of $\cup S$. This is already Condition (a) above. Condition (b) holds by Lemma 3. Therefore, $\cup R$ is a maxcon of $(R|S \setminus R)$. What remains to be proved to establish that this partition is stable is that $\cup R$ induces it. The first class of the partition induced by $\cup R$ contains all sources in S that are consistent with $\cup R$, and the second class the sources

that are inconsistent with $\cup R$. These are respectively R and $S \setminus R$ because $\cup R$ is consistent and maximally so in $\cup S$. \square

The maxcons of the stable partitions can be taken as the result of merging. The above theorem proves that in the case of singleton sources and two-class partitions, this is the same as selecting the plain maxcons of the set of formulae provided by the sources.

4. Two-Class Partitions of Non-singleton Sources

The previous section analyzes the simplified case where each source provides exactly one piece of information. In the example of the mayor and the landslide, this is like each expert made exactly one statement in their report. This is not the case, as Bialelli and Connors made two statements each. Again, the mayor may choose any reliability ordering since no one is known. Yet, when integrating formulae using the chosen order the result should not be inconsistent with the ordering itself; otherwise, the mayor would get very bad press for not being able to follow his own opinion.

Theorem 2 shows that no selection stems from assuming and then checking the reliability of single-formulae sources: the plain maxcons are generated exactly by the stable partitions. This is unsurprising, since the reliability of a formula is the same as the reliability of its source, and is only affected by the other formulae in the same source. This is the case considered in this section: sources with multiple formulae.

Definition 8 *A set of formulae M induces on a set of sources S the ordered partition $(R|S \setminus R)$ where*

$$R = \{S_i \mid M \not\models A_i \text{ for all } A_i \in S_i\}$$

The partition is on the sources, not on formulae. As a result, even if the same formula is provided by two sources, these may be in different classes.

Definition 9 *An ordered partition of the sources is **weakly stable** if it has a maxcon M that induces the partition itself.*

This definition requires only that one of the maxcons of the partition induces the partition itself. But a partition may have multiple maxcons; the others may induce different partitions. Requiring the condition to hold for all of them will be shown to have different consequences.

Definition 10 *An ordered partition $(R|S \setminus R)$ of the sources is **strongly stable** if all its maxcons induce the partition itself.*

As an example, a partition P may have two maxcons M' and M'' . Each of these maxcons induces a partition: M' induces P' and M'' induces P'' . If $P' = P$ then P is weakly stable since it is obtained by going from it to a maxcon and from the maxcon to a partition. If $P'' \neq P$ then P is not strongly stable, since a different partition can also be reached in the same way. The following Theorem 4 shows such a partition that is weakly but not strongly stable.

A different way to restrict the maxcons is to still consider all weakly stable partitions, but only their maxcons that induce the partition itself. It will however be shown that the other maxcons induce weakly stable partitions. In other words, a maxcon that does not induce the partition itself induces another partition that is weakly stable.

Yet another restriction is to start from the partition having all sources in its first class. This realizes the assumption that all sources are initially compared the same. Following the maxcons and the induced partitions may then lead to some weakly or strongly stable partitions. This may seem a restriction since only partitions reachable from the all-equal one are considered. As shown below, it is not.

The results proved in this section for non-singleton sources and two-class partitions are:

- restricting to partitions reachable from the all-equal one does not eliminate any weakly or stable partition;
- some weakly stable partitions are not strongly stable;
- a weakly or strongly stable partition may have more than one maxcon;
- the maxcons of the weakly stable partitions are the plain maxcons;
- a weakly stable partition may have a maxcon that induces a different partition; however, the first class of the latter is larger than that of the former;
- the maximal weakly stable partition w.r.t. containment of their first classes are the strongly stable partitions;
- the maxcons of the strongly stable partitions are also the maxcons of the specific partition $(\wedge S_1, \dots, \wedge S_n | S_1 \cup \dots \cup S_n)$, where $\wedge S_i$ is the conjunction of the formulae in S_i and $S_1 \cup \dots \cup S_n$ is the set of all formulae in these sources.

Considering only partitions reachable from the partition that has all sources in its first class is not a restriction at all, even if reachability is limited to a single step.

Theorem 3 *If a partition $(R|S \setminus R)$ is weakly or strongly stable then some maxcon M of the partition (S) induces $(R|S \setminus R)$.*

Proof. Since the partition is stable, it has a maxcon M that induces itself. By definition, M is also a plain maxcon. The partition (S) has exactly all plain maxcons, including M . \square

Every strongly stable partition is also weakly stable, as an immediate consequence of the definition. The converse does not always hold: a weakly stable partition may have a maxcon that induces a partition that is different from the original one.

Theorem 4 *There exists a set of sources whose weakly stable partitions are not all strongly stable.*

Proof. The sources are $\{A\}$, $\{B, C\}$ and $\{D\}$, the plain maxcons $\{A, B\}$, $\{B, C\}$, $\{A, D\}$. This is a correct description of a merging scenario because sources are identified with the set of formulae they provide, and formulae can be specified by giving their maxcons thanks




Figure 4: The formulae in the proof of Theorem 4: A is consistent with D and B , which is consistent with C .

to Lemma 5. Figure 4 show the situation in the space of models: A shares models with D and with B , which in turn shares other models with C .

All maxcons of the partition $(\{A\}|\{B, C\}\{D\})$ contain A by Lemma 2. Of the formulae of the second class, only B and D are consistent with A , leading to the maxcons $\{A, B\}$ and $\{A, D\}$. Their induced partitions are:

$\{A, B\}$: the only source that contains only formulae consistent with this maxcon is $\{A\}$; as a result, the induced partition is $(\{A\}|\{B, C\}\{D\})$ itself, showing that this partition is weakly stable;

$\{A, D\}$: this maxcon is consistent with the two sources $\{A\}$ and $\{D\}$; the induced partition is therefore $(\{A\}\{D\}|\{B, C\})$, which is not the same as the partition assumed initially.

The first point shows that $\{A\}|\{B, C\}\{D\}$ is weakly stable, since it has a maxcon that induces the partition itself. The second shows that it is not strongly stable, since the other maxcon induces a different partition. \square

This theorem proves that a partition may be weakly but not strongly stable. This is only possible if the partition has at least two maxcons: at least one induces the partition itself and at least one does not. A related question is whether all partitions having multiple maxcons are in this condition.

A minor change in the sources of the proof confutes this supposition. Adding a single other formula E that is only consistent with D generates a strongly stable partition with two maxcons, as shown in Figure 5. The partition $(\{A\}|\{B, C\}\{D, E\})$ has two maxcons: $\{A, B\}$ and $\{A, D\}$. The source $\{A\}$ is the only one having only formulae consistent with the first maxcon, leading to the partition itself. The same for the second maxcon. This is a strongly stable partition that has two maxcons.

As in the case of singleton sources, the maxcons of the weakly stable partitions are exactly all the plain maxcons of the set of sources. This is proved from a property of the plain maxcons.

Lemma 7 *If M is a plain maxcon, its induced partition has M as a maxcon.*

Proof. Let $(R|S\setminus R)$ be the partition induced by M . By definition, all formulae in R are consistent with M . If some of these formulae were not in M , then M would not be a plain




Figure 5: A scenario in which a strongly stable partition has two maxcons: the partition is $(\{A\}|\{B, C\}\{D, E\})$ and the maxcons $\{A, B\}$ and $\{A, D\}$.

maxcon of S . This proves that M contains all formulae in R . By Lemma 3, M is a maxcon of $(R|S \setminus R)$. \square

This lemma shows that starting from a plain maxcon and moving to its induced partition and then to its maxcons, among them is the original maxcon. The partition may also have other maxcons, as shown in the proof of Theorem 4: the maxcon $\{A, B\}$ induces the partition $(\{A\}|\{B, C\}\{D\})$, which has two maxcons: $\{A, B\}$ itself and $\{A, D\}$.

This lemma has two side consequences.

The first is related to an additional constraint that was considered above: instead of accepting all maxcons of all weakly stable partitions, restrict to those generating the partition. In other words, if a weakly stable partition $(P_1|P_2)$ has the maxcon M , but this maxcon induces a different partition, exclude M from consideration. This restriction is useless: even if M induces a different partition $(P'_1|P'_2)$, one of the maxcons of this partition is M itself.

The second consequence is that every set of sources has at least a weakly stable partition. Indeed, every set has at least a maxcon M , and this maxcon induces a partition that has M as a maxcon. That partition is therefore weakly stable.

The converse of the lemma holds by definition: the maxcons of a partition are some of the plain maxcons.

Corollary 1 *The maxcons of weakly stable partitions are the plain maxcons.*

This corollary allows for a method for finding the weakly stable partitions and their maxcons at the same time. The most obvious way for obtaining a plain maxcon is to start from the empty set and then iteratively adding a formula if consistent with the set. When formulae are provided in a set of sources, this method can be slightly simplified because M has to be consistent with all formulae in the first class of some partition.

Algorithm 1 (maxcon and induced partition of two classes)

Input: a set of sources $S = \{S_1, \dots, S_n\}$

Output1: a weakly stable partition $(R|S \setminus R)$

Output2: a maxcon M inducing $(R|S \setminus R)$ and induced by $(R|S \setminus R)$

1. $M = \emptyset$
2. $R = \emptyset$
3. for each source $S_i \in S$:
 - if $M \cup (\cup S_i)$ is consistent:
 - i. $M = M \cup (\cup S_i)$
 - ii. $R = R \cup \{S_i\}$ // remark: here R changes
4. output the weakly stable partition $(R|S \setminus R)$
5. for each $S_j \in S \setminus R$ // remark: use R updated above
 - for each $F \in S_j$
 - if $S_j \setminus \{F\} \not\subseteq M$ and $M \cup \{F\}$ is consistent then
 - $M = M \cup \{F\}$
6. output maxcon M

The algorithm first finds a stable partition by iteratively adding entire sources to its maxcon in construction M ; once the partition is output in Step 4, the process of addition continue on single formulae from the remaining sources $S \setminus R$.

The test $S_i \setminus \{F_i\} \not\subseteq M$ is not necessary but may save some consistency checks, especially for sources of low cardinality: since S_i is not in the first class of the partition induced by M , not all its formulae can be in M . Therefore, if all of them but F are in M , then $M \cup \{F\}$ is inconsistent.

Theorem 5 *Algorithm 1 outputs a weakly stable partition and one of its maxcons inducing the partition itself.*

Proof. The algorithm always terminates because its two loops are over finite sets: S and $S \setminus R$. The first loop has the following invariants:

1. $M = \cup R$; this is because both R and M start empty, and when a source is added to M its formulae are added to M ;
2. M is consistent; this is because M is only added S_i if $M \cup (\cup S_i)$ is consistent;
3. if R does not contain S_i after S_i has been analyzed in the loop, then $\cup R \models \neg(\wedge S_i)$; the contrary implies the consistency of $(\cup R) \cup S_i$, which is now proved to contradict the assumption; let M' and R' be the state of M and R during the iteration of S_i ; it holds $R' \subseteq R$ since sources are never removed from R ; it holds $M' = \cup R'$ because of the first invariant; therefore, $M' \cup S_i$ is consistent and S_i is inserted into R' at this iteration; since $R' \subseteq R$, this contradicts the assumption that R does not contain S_i .

At the end of the first loop all sources S_i are analyzed. The second invariant therefore becomes: R contains all S_i such that $(\cup R) \cup S_i$ is consistent. The first invariant is still $M = \cup R$. This implies that M contains all formulae in all sources S_i such that $M \cup S_i$ is consistent.

The partition $(R|S \setminus R)$ is output at this point, but R is not changed afterwards. Therefore, if this partition is weakly stable after the second loop, it is also weakly stable at this point. The proof of weak stability of this partition can therefore be delayed after the second loop.

The invariants of the second loop are: first, M is consistent; second, if F is a formula in R or in a source S_i already analyzed in the second loop, then $F \notin M$ if and only if $M \cup \{F\}$ is inconsistent. The first invariant holds because M is only added a formula if it is consistent with it. The second holds because of what holds at the end of the first loop and because a formula is not added to M only if inconsistent with it.

At the end of the second loop all formulae F have been considered; therefore, M contains all formulae F it is consistent with. Since M is also consistent, it is a plain maxcon.

Since formulae are never removed from M , this set still contains all formulae in R . As a result, it is a maxcon of $R \cup (S \setminus R)$ containing all formulae in R , which implies it is a maxcon of $(R|S \setminus R)$. What remains to be proved is that M induces this very partition. Since the considered partitions comprise two classes, this is the same as R being the first class of the partition induced by M . Since M is consistent and contains all formulae in R , all sources in R are in the first class of the partition induced by M . If a source $S_i \notin R$ were in the first class, then $M \cup S_i$ would be consistent, implying the consistency of $(\cup R) \cup S_i$ since M contains all formulae in R . This contradicts the invariant at the end of the first loop. \square

A weakly stable partition $(R|S \setminus R)$ induces at least a maxcon M whose induced partition is $(R|S \setminus R)$. However, the same partition may also induce other maxcons, such as M' and M'' . This raises two questions: Which properties these other maxcons have? Which partitions they induce?

Theorem 6 *If $(R|S \setminus R)$ is weakly stable, then all its maxcons include $\cup R$ and induce a partition $(R'|S \setminus R')$ such that $R \subseteq R'$.*

Proof. Since the partition is weakly stable, it has a maxcon M that induces the same partition $(R|S \setminus R)$. Since M induces $(R|S \setminus R)$, all formulae in R are consistent with M . By Lemma 1, since M is a maxcon then it contains them as well. This proves that R is contained in M ; therefore, $\cup R$ is consistent; by Lemma 2, all maxcons of $(R|S \setminus R)$ include $\cup R$ and all their induced partitions have all of R in their first class. \square

If a weakly stable partition $(R|S \setminus R)$ has a maxcon M that induces a partition $(R'|S \setminus R')$, then either $R = R'$, which implies that the two partitions are the same, or $R \subset R'$. If a partition has a maxcon that induces a different partition, the first class of the latter is strictly larger than that of the former.

In other words, when moving from partition to maxcon to a different partition, the first class strictly increases. This forbids a path partition—maxcon—partition—... to form a cycle, unless all partitions in it are the same. Indeed, in such a cycle it would be $R \subset R' \subset R'' \dots \subset R$, which is impossible.

From a different angle, since the first class strictly increases, at some point a maximum is hit: a partition only has maxcons that induce the partition itself.

Theorem 7 *The maximal weakly stable partitions w.r.t. containment of their first class are the strongly stable partition.*

Proof. Let $(R|S\setminus R)$ be a weakly stable partition that is maximal on set-containment of its first class. It is not strongly stable only if one of its maxcons M induces a partition $(R'|S\setminus R')$ with $R \neq R'$. Theorem 6 implies $R \subseteq R'$, which implies $R \subset R'$ in this case. Since $(R'|S\setminus R')$ is induced by M , it has M as a maxcon by Lemma 7. Therefore, this partition is weakly stable and has its first class strictly containing R , contrary to the maximality of $(R|S\setminus R)$.

The converse is proved by assuming that $(R|S\setminus R)$ is strongly stable, and proving that it is a maximal weakly stable partition. That it is weakly stable is an immediate consequence of the definitions of stability. Remain to prove that it is maximal. Let $(R'|S\setminus R')$ be another weakly stable partition with $R \subset R'$. By definition of weak stability, $(R'|S\setminus R')$ has a maxcon M whose induced partition is $(R'|S\setminus R')$. By Theorem 6, M contains $\cup R'$, which in turn contains $\cup R$. Since every maxcon is also a plain maxcon, and M contains $\cup R$, by Lemma 3 M is a maxcon of $(R|S\setminus R)$. As a result, M is a maxcon of $(R|S\setminus R)$ but its induced partition is $(R'|S\setminus R')$, contradicting the assumption of strong stability. \square

This proves the existence of strongly stable partitions for all set of sources.

Theorem 8 *Every set of sources has at least a strongly stable partition.*

Proof. Lemma 7 tells that the partitions induced by plain maxcons are weakly stable, and therefore proves the existence of weakly stable partitions for all set of sources. By Theorem 6, if such a partition is not strongly stable then it has a maxcon that induces a partition with a strictly larger first class. By Lemma 7 and the fact that maxcons of partitions are also plain maxcons, this partition is weakly stable. Since the sources are finitely large, moving from a partition to another must eventually hit a maximum. By Theorem 7, these maximal partitions are strongly stable. \square

The proof also shows a way to find the strongly stable partitions from weakly stable partitions: following the maxcons and their induced partitions.

Algorithm 2 (Weakly and strongly stable partitions of two classes)

Input: a set of sources $S = \{S_1, \dots, S_n\}$

Output1: a weakly stable partition

Output2: a strongly stable partition

1. set P to be an arbitrary partition
2. find a maxcon M of P // see below
3. determine the partition P' that M induces
4. output P' as a weakly stable partition (Lemma 7)

5. for each maxcon M' of P' :

if M' induces a partition $P'' \neq P'$ then:

i. set $P' = P''$

ii. go to Step 5

6. output P' as strongly stable partition (Theorem 7)

Step 2 can be done in a number of ways. The maxcon can be determined by setting $M = \emptyset$ and iteratively adding the formulae in the first class of P that are consistent with M , then from the second, etc. Alternatively, Algorithm 1 may replace Steps 1-3, since it produces both M and one of its induced partitions P' at the same time; P is in this case unnecessary, since it is not used in the rest of the algorithm.

Theorem 9 *Algorithm 2 outputs a weakly and a strongly stable partition.*

Proof. Every maxcon M of a partition P is also a plain maxcon by definition, and the partition induced by M has M as a maxcon thanks to Lemma 7. This proves that Step 4 outputs a weakly stable partition.

This partition may also be strongly stable or not. If it is, none of its maxsets induce another partition; therefore, the condition in Step 5 is always false; the partition is not changed, and is output as is in Step 6.

If this partition P' is not strongly stable, the condition in Step 5 is true for some maxset M' of P' . Therefore, the partition P'' induced by M' replaces P' . By Theorem 6, P'' has a strictly larger first class than the previous P' . The check is performed again. Since the formulae are finite and the first class strictly increases at each step, at some point this loop terminates. This means that the condition of the loop is false: no maxcon M' of P' induces a partition P'' different from P' . This is the definition of strongly stability. \square

The strongly stable partitions can also be found employing a specific partition. Let $\wedge S_i$ be the conjunction of the formulae in the source S_i . The partition is $(\{\wedge S_1, \dots, \wedge S_n\} | S_1 \cup \dots \cup S_n)$. For example, the sources $S_1 = \{A, B\}$, $S_2 = \{C, D, E\}$ and $S_3 = \{F\}$ are turned into the partition $(\{A \wedge B, C \wedge D \wedge E, F\} | \{A, B, C, D, E, F\})$. The following theorem proves that the maxcons of this partition are the maxcons of the strongly stable partitions of the three original sources S_1 , S_2 and S_3 .

Theorem 10 *For every maxcon M of a strongly stable partition of the sources S_1, \dots, S_n there exists a maxcon M' of the partition $(\{\wedge S_1, \dots, \wedge S_n\} | S_1 \cup \dots \cup S_n)$ such that $M = M' \cap (S_1 \cup \dots \cup S_n)$, and vice versa.*

Proof. Let $S = \{S_1, \dots, S_n\}$ and $S^\wedge = \{\wedge S_1, \dots, \wedge S_n\}$. The partition in the statement of the theorem is $(S^\wedge | \cup S)$. Given $M \subseteq \cup S$, let $a(M)$ be the set defined by:

$$a(M) = M \cup \{\wedge S_i \in S^\wedge \mid M \cup \{S_i\} \neq \perp\}$$

Let M be a maxcon of a strongly stable partition $(R|S\setminus R)$ and prove that $a(M)$ is a maxcon of $(S^\wedge| \cup S)$.

If M is consistent with $\wedge S_i$, every formula in S_i is consistent with M . Therefore, it is in M since M is also a maxcon of S . This implies that all formulae $\wedge S_i$ consistent with M are also entailed by M . Therefore, M and $a(M)$ are equivalent.

Since M is a maxcon of $(R|S\setminus R)$, it is a maxcon of its flattening $(\cup R| \cup (S\setminus R))$. By definition, it is also a maxcon of the union of the classes $(\cup R) \cup (\cup((S\setminus R))) = \cup S$. Since $a(M)$ contains all elements of S^\wedge that are consistent with M , it is a maxcon of $S^\wedge \cup (\cup S)$. In other words, it is a plain maxcon of $(S^\wedge| \cup S)$. This is the first part of the definition of $a(M)$ being a maxcon of $(S^\wedge| \cup S)$; the other that is still to be proved is that $a(M) \cap S^\wedge$ is a maxcon of S^\wedge .

The converse of this condition is possible if a source is consistent with $a(M) \cap S^\wedge$ but is not contained in it. For the sake of clarity, the sources can be re-indexed so that $a(M) \cap S^\wedge = \{\wedge S_1, \dots, \wedge S_i\}$ and the source that is consistent but not contained in this set is S_{i+1} .

Using such indexes, the partition induced by M is $(S_1 \dots S_i | S_{i+1} \dots S_n)$; this is the same as $(R|S\setminus R)$ thanks to strong stability. Since M is a maxcon, it is consistent. Therefore, its equivalent set $a(M)$ and its subset $a(M) \cap S^\wedge$ are consistent as well. The latter is also consistent with $\wedge S_{i+1}$ by assumption; therefore, all formulae in S_{i+1} can be added to it while retaining consistency. Iteratively adding formulae from $S_{i+2} \dots S_n$ produces a maxcon M' of the partition $(R|S\setminus R)$. The partition induced by M' has S_{i+1} in the first class and is therefore different from $(R|S\setminus R)$, contrary to the assumption of strong stability.

The second part of the proof is to show that every maxcon of $(S^\wedge| \cup S)$ is also a maxcon of a strongly stable partition of S after the formulae $\wedge S_i$ are removed from it. This is achieved by showing that the partition induced by the maxcon is strongly stable.

Let M be a maxcon of $(S^\wedge| \cup S)$. Let $M' = M \cap S^\wedge$ and $M'' = M \cap (\cup S)$. Since M is maximally consistent, $\wedge S_i \in M'$, $M'' \cup S_i \not\equiv \perp$ and $S_i \subseteq M''$ are equivalent to each other. This proves that M and M'' are equivalent.

The partition induced by M'' on S is $(R|S\setminus R)$, where $R = \{S_i \mid M'' \cup S_i \not\equiv \perp\}$, which can be rewritten as $R = \{S_i \mid \wedge S_i \in M'\}$. Since M' is consistent, so is R .

This partition is not strongly stable only if it has a maxcon M''' that induces a different partition. By Lemma 2, since R is consistent M''' contains all its formulae. Therefore, the partition induced by M''' has all of R in the first class. This partition can be different from $(R|S\setminus R)$ only if some $S_i \in S\setminus R$ is in its first class, which requires it to be consistent with M''' , and therefore with R . This implies that the formulae of S^\wedge that are consistent with M''' form a strict superset of those consistent with M' , contrary to the assumption that M is a maxcon. \square

5. Three-Class Partitions

The mayor classed the experts on landslides as fully reliable or not: Akai and Bialetti are reliable, Connors less so. Such a division may be too strict, on a field that is so complex that even experts make mistakes. For example, some experts may be totally reliable (for example, Akai is an expert on both geology in general and that area in particular), some partly reliable (Bialetti authored many studies on landslides, but does not have any specific knowledge of

the geology of that particular zone), others even less reliable (Connors' preparation might be outdated).

If such a three-classes division were known, merging could be performed by taking the statements in order of decreasing reliability. But the mayor has no idea of the expert's reliability. He can easily make up excuses for justifying any order of reliability he chooses, but the result of merging should then be consistent with that order. This section analyzes this problem when sources are framed into three reliability classes.

The same fixpoint definition of the previous sections can still be used: a three-class ordered partition is assumed, and its maxcons are found; each induces a partition, where the first class is defined as before, but the second comprises only the sources that provide some formulae that are consistent with the maxcon. The third class contains all other sources, whose formulae are all inconsistent with the maxcon. According to the maxcon, the sources in the first group are more reliable of those in the second, in turn more reliable than those in the third.

Weak and strong stability are defined as before: a partition is weakly or strongly stable if one or all of its maxcons induce the partition itself.

The only difference between the previous cases is the definition of induced maxcon, which is now composed of three classes. The rest of the framework is as before.

Definition 11 *A set of formulae M induces on the set of sources $S = \{S_1, \dots, S_n\}$ the partition of three classes $(R|P|U)$ where*

$$\begin{aligned} R &= \{S_i \mid M \not\models \neg A \text{ for all } A \in S_i\} \\ P &= \{S_i \mid M \not\models \neg A \text{ for some but not all } A \in S_i\} \\ U &= \{S_i \mid M \models \neg A \text{ for all } A \in S_i\} \end{aligned}$$

If the mayor takes a decision based on the integrated information M , it follows that according to the mayor's information the experts in R are more reliable than those in P , in turns more reliable than those in U . Definition 11 can be used in the same fixpoint condition of the previous sections, so that a three-class partition is stable if its maxcons induce the partition itself. Depending on whether all or some of these maxcons do, stability is strong or weak.

The results in this section concern ordered partitions of three classes. Some relate partitions of two classes and partitions of three. To avoid confusion, these are called *bipartitions* and *tripartitions*, respectively. Given the number of theorems in this section, a summary is in order:

- the tripartition induced by a plain maxcon has it as a maxcon;
- weakly stable bipartitions correspond to weakly stable tripartitions and vice versa;
- strongly stable tripartitions may not correspond to strongly stable bipartitions;
- some weakly stable bipartitions can be turned into strongly stable tripartitions, but not all;

- when moving from tripartition to maxcon to tripartition, the first class increases or stays the same, like it does for bipartitions; the same holds neither for the second class nor its union with the first; the third class of the resulting tripartition may be incomparable with the original;
- cycles of weakly stable tripartitions are possible; there are even cases where this holds for all maxcons;
- not all sets of sources have strongly stable tripartitions;
- the relation linking weakly stable tripartitions and maxcons is not transitive nor symmetric; the latter holds even if the tripartitions have the same R ; furthermore, the relation can be made to include an arbitrary graph of tripartitions.

The first result carries Lemma 7 over from bipartitions to tripartitions.

Theorem 11 *If M is a plain maxcon of S , it is also a maxcon of its induced tripartition.*

Proof. Let $(R|P|U)$ be the tripartition induced by M . By definition, M is a maxcon of this tripartition if it is a maxcon of $\cup(R \cup P \cup U)$, its intersection with $\cup(R \cup P)$ a maxcon of $\cup(R \cup P)$ and its intersection with $\cup R$ a maxcon of $\cup R$. The first condition holds by assumption, as M is a plain maxcon. For the other two conditions, since M is consistent only maximal consistency within $\cup(R \cup P)$ and $\cup R$ is left to prove.

Regarding $\cup(R \cup P)$, by definition of induced tripartition $U = \{S_i \mid \forall F \in S_i . M \models \neg F\}$. Since M is consistent, it does not contain any formula in U and therefore only comprises formulae of the other two classes: $M \cap \cup(R \cup P) = M$. If there exists M' such that $M \cap \cup(R \cup P) \subset M' \cap \cup(R \cup P)$ then $M \subset M'$; therefore, M' is inconsistent because M is a plain maxcon of S .

Regarding $\cup R$, by definition of induced tripartition $R = \{S_i \mid \forall F \in S_i . M \not\models \neg F\}$. Since M is a plain maxcon, if it is consistent with F then it includes it. As a result, M contains all formulae in R . Therefore, $M \cap \cup R$ is a maxcon of $\cup R$ simply because it contains the whole of it. \square

Plain maxcons are also maxcons of a tripartition if they obey a simple condition.

Lemma 8 *If a plain maxcon of $\cup(R \cup P \cup U)$ contains all formulae in R and none in U , it is a maxcon of the tripartition $(R|P|U)$.*

Proof. The claim requires $M \cap (\cup R)$ to be a maxcon of $\cup R$ and $M \cap (\cup(R \cup P))$ to be a maxcon of $\cup(R \cup P)$. For the first condition: $M \cap (\cup R)$ is consistent because M is consistent (being a plain maxcon); it is maximal because it contains all formulae in R .

The second condition also holds. Since M is a plain maxcon of $\cup(R \cup P \cup U)$, it holds $M \subseteq \cup(R \cup P \cup U)$. Since $M \cap (\cup U) = \emptyset$ also holds (by assumption), $M \subseteq \cup(R \cup P)$ follows. Therefore, $M \cap (\cup(R \cup P)) = M$. Since M is maximally consistent within $\cup(R \cup P \cup U)$, it is also maximally consistent within $\cup(R \cup P)$. \square

The converse is not always the case. A maxcon may contain formulae in the third class, as for example $\{x, y\}$ is a maxcon of $(\{x\}|\{\neg x\}|\{y\})$. This is exactly one of the cases where

the approach of assuming a partition and then checking it via its maxcons yields its fruit: the reliability ordering initially assumed reckoned $\{\neg x\}$ more reliable than $\{y\}$, but merging under this ordering produces a result that is inconsistent with the more reliable source and consistent with the less. Disproving the relative reliability of sources this way is exactly one of the aims of the method of assuming a partition, determining a maxcon and then checking the original partition.

Lemma 9 *Some weakly stable tripartitions have maxcons that contain formulae in their third class.*

Proof. The maxcons of the tripartition $(\{x \vee y\}|\{\neg x \wedge \neg z, \neg y\}|\{z\})$ are $\{x \vee y, \neg x \wedge \neg z\}$ and $\{x \vee y, \neg y, z\}$. The first induces the tripartition itself, proving that weakly stable. The second contains the formula in the third class, and therefore does not induce the same tripartition. \square

As far as weak stability is concerned, using two or three classes is the same: weakly stable bipartitions and weakly stable tripartitions correspond to each other.

Theorem 12 *For each weakly stable bipartition $(R|S \setminus R)$ there exists a weakly stable tripartition $(R|P|U)$ where $P \cup U = S \setminus R$, and vice versa.*

Proof. If $(R|S \setminus R)$ is weakly stable, it has a maxcon M that induces the bipartition itself. The tripartition induced by M has the same first class of the bipartition because the definition of the first class is the same for bipartitions and tripartitions. Let $(R|P|U)$ be this tripartition. Since M is also a plain maxcon, $(R|P|U)$ has M as a maxcon by Lemma 11. Since M induces $(R|P|U)$, this tripartition is weakly stable.

The other direction is proved in the same way. Let $(R|P|U)$ be a weakly stable tripartition, and M be a maxcon that induces the tripartition itself. The tripartition induced by M is $(R|P|U)$; since the first class of induced bipartitions and tripartitions are the same, the bipartition induced by M is $(R|S \setminus R)$. By Lemma 7, this bipartition has M as a maxcon, which proves it to be weakly stable. \square

While weakly stable bipartitions and tripartitions are the same apart from the split of the second class, strongly stable bipartitions and tripartitions exhibit quite a difference. Indeed, a bipartition may not be strongly stable because of a maxcon that induces a different bipartition; splitting the second class in two (which produces a tripartition) may block such a maxcon.

Theorem 13 *There exists a strongly stable tripartition $(R|P|U)$ such that $(R|P \cup U)$ is weakly but not strongly stable.*

Proof. The claim is proved on the same sources as in Theorem 4: $S_1 = \{A\}$, $S_2 = \{B, C\}$, and $S_3 = \{D\}$, with maxcons $\{A, B\}$, $\{B, C\}$ and $\{A, D\}$. These formulae are depicted in the space of models in Figure 6.

As already proved, the bipartition $(\{A\}|\{B, C\}, \{D\})$ is weakly but not strongly stable because of the maxcon $\{A, D\}$, which induces the different partition $(\{A\}, \{D\}|\{B, C\})$. As a result, its other maxcon $\{A, B\}$ is not a maxcon of any strongly stable partition, since the partition it induces is $(\{A\}|\{B, C\}, \{D\})$, which is not strongly stable.




Figure 6: These formulae show a strongly stable tripartition which, when turned into a bipartition, is only weakly stable.




Figure 7: The weakly stable tripartition $(\{A\}|\{B, C\}, \{D, E\})$ does not correspond to a strongly stable bipartition.

The bipartition $(\{A\}|\{B, C\}, \{D\})$ is related to the tripartition $(\{A\}|\{B, C\}|\{D\})$ as specified by the claim: the set of all sources and the first class are the same. The tripartition $(\{A\}|\{B, C\}|\{D\})$ is strongly stable, as its only maxcon $\{A, B\}$ induces the tripartition itself: it is consistent with A , it is consistent with one but not all formulae in $\{B, C\}$, and is inconsistent with D . This proves that a bipartition that is not strongly stable can be turned into a strongly stable tripartition. \square

While some bipartitions can be turned from weakly to strongly stable by splitting their second class, this is not the case for all of them.

Theorem 14 *There exists a weakly stable tripartition $(R|P|U)$ such that $(R|P \cup U)$ is weakly but not strongly stable.*

Proof. The sources are $\{A\}$, $\{B, C\}$, and $\{D, E\}$; their plain maxcons are $\{A, B, E\}$, $\{B, C\}$ and $\{A, D, E\}$. The situation in the space of models is depicted in Figure 7.

The bipartition $(\{A\}|\{B, C\}, \{D, E\})$ is weakly but not strongly stable. Indeed, it has two maxcons: $\{A, B, E\}$ and $\{A, D, E\}$. The first induces the bipartition itself, proving it weakly stable. The second induces the different bipartition $(\{A\}, \{D, E\}|\{B, C\})$, proving that stability is not strong.




Figure 8: The formulae in the proof of Theorem 15.

The maxcon $\{A, B, E\}$ induces the tripartition $(\{A\}|\{B, C\}, \{D, E\}|)$. This tripartition has two maxcons: $\{A, B, E\}$ and $\{A, D, E\}$. The second maxcon $\{A, D, E\}$ induces the different tripartition $(\{A\}, \{D, E\}|\{B, C\})$ (the source $\{B, C\}$ is in the third class since both B and C are inconsistent with $A \wedge D \wedge E$). As a result, the tripartition is weakly but not strongly stable. \square

Finally, some strongly stable bipartitions cannot be turned into a strongly stable tripartition, no matter how the second class is split. It will indeed be proved that some sets of sources do not have any strongly stable tripartition, while strongly stable bipartitions are guaranteed to exist as a consequence of Theorem 7 or of Theorem 10. Such a set is provided in the following Corollary 2.

Theorem 6 proves a number of results about the first class of weakly stable bipartitions: they have a consistent first class, their maxcons contain all of it, and the partitions they induce have a larger or equal first class. All of this extend to tripartitions because the proof was based on the first class, which is defined in the same way. However, nothing like this holds for the second and the third class of a tripartition.

Theorem 15 *A weakly stable tripartition $(R|P|U)$ of some sources has a maxcon that induces a partition $(R'|P'|U')$ such that $P \not\subseteq P'$, $R \cup P \not\subseteq R' \cup P'$, $U \not\subseteq U'$ and $U' \not\subseteq U$.*

Proof. The claim is proved using the sources $\{A, B, C\}$, $\{D, E\}$ and $\{F\}$ with the plain maxcons $\{A, B, C\}$, $\{B, C, D\}$ and $\{D, E, F\}$. Figure 8 shows such formulae in the space of models.

A weakly stable tripartition of these sources is $(|\{A, B, C\}, \{D, E\}|\{F\})$, since its maxcon $\{B, C, D\}$ induces the partition itself.

This tripartition also has the two maxcons $\{A, B, C\}$ and $\{D, E, F\}$. The first maxcon $\{A, B, C\}$ induces the partition $(\{A, B, C\}|\{D, E\}\{F\})$, which shows that the second class does not necessarily increase, neither does its union with the first. The second maxcon $\{D, E, F\}$ induces $(\{D, E\}, \{F\}|\{A, B, C\})$, whose third class is incomparable with that of the original tripartition. \square

Incidentally, the set of sources has a strongly stable tripartition. Indeed, the only maxcon of the tripartition $(\{A, B, C\}|\{D, E\}, \{F\})$ is $\{A, B, C\}$, which induces the partition itself.




Figure 9: The formulae used to prove that switching from a weakly stable tripartition to the tripartition induced by one of its maxcons the second class may decrease.

In this and in the previous proof, going from a tripartition to a maxcon and to its induced tripartition reduces the second class. This is however easy to show not to always be the case, with the same sources but $\{F, G\}$ in place of $\{F\}$, where G is only consistent with F . Thanks to this change, the maxcon $\{D, E, F\}$ induces the tripartition $(\{D, E\}|\{F, G\}|\{A, B, C\})$.

Another common feature of the above proofs is that the tripartitions that are weakly but not strongly stable have a maxcon whose induced tripartition has a second class that is strictly larger than the original tripartition. This is however not always the case.

Theorem 16 *There exists a tripartition with two maxcons such that the first induces the tripartition itself and the second induces a tripartition with a smaller second class.*

Proof. The proof is based on the tripartition $(\{A\}|\{B, C, D\}, \{E, F\}|)$ whose plain maxcons are $\{A, B, C\}$, $\{A, B, E\}$, $\{B, C, D\}$ and $\{E, F\}$, as depicted in Figure 9.

One of the maxcons of the tripartition $(\{A\}|\{B, C, D\}, \{E, F\}|)$ is $\{A, B, E\}$, which induces the tripartition itself. This is the first part of the claim. The only other maxcon of the tripartition is $\{A, B, C\}$, which induces $(\{A\}|\{B, C, D\}|\{E, F\})$. This tripartition has the same first class of the original one, but a smaller second class. \square

A strongly stable bipartition can always be found by starting from a maxcon and then following the induced bipartition and then its maxcons, or starting from a weakly stable bipartition and following a similar path. Nothing like this can be done on tripartitions: one may start from a weakly stable tripartition, follow the maxcons and then the induced tripartitions and so only to end up in the original tripartition. This can be the case even for all maxcons of a tripartition, as the following example shows.

Lemma 10 *There exist two different tripartitions T and T' that have the same first class and the same two maxcons M and M' , where M induces T and M' induces T' .*

Proof. The sources are $\{A\}$, $\{B, C\}$, $\{D, E\}$, $\{F, G\}$ and their plain maxcons are $\{B, C\}$, $\{A, B, D\}$, $\{D, E\}$, $\{A, E, F\}$ and $\{F, G\}$. They are depicted in Figure 10.




Figure 10: The formulae in the proof of Theorem 10.

The tripartitions and the maxcons of the claim are:

$$\begin{aligned}
 T &= (\{A\}|\{B, C\}, \{D, E\}|\{F, G\}) \\
 T' &= (\{A\}|\{F, G\}, \{D, E\}|\{B, C\}) \\
 M &= \{A, B, D\} \\
 M' &= \{A, E, F\}
 \end{aligned}$$

Both tripartitions have A only in their first class. All their maxcons therefore contain A . The only plain maxcons containing A are M and M' . Both are maxcons of the tripartitions, since they cannot be consistently extended by a formula of the second class even removing the formulae of the third (F for the first tripartition and B for the second). That M induces T and M' induces T' is a consequence of the definition. \square

A first consequence of this lemma is the existence of strongly stable bipartitions that cannot be turned into strongly stable tripartitions by splitting their second class.

Corollary 2 *There exists a strongly stable bipartition $(R|S \setminus R)$ such that $(R|P|U)$ is not strongly stable for any P and U such that $R \cup P \cup U = S$.*

Proof. The tripartitions T and T' of the previous theorem are weakly but not strongly stable tripartitions. Since their first classes coincide, they are associated with the same bipartition, which is strongly stable because its maxcons M and M' induce the tripartition, and therefore induce the same bipartition. \square

Lemma 10 also proves that a strongly stable tripartition cannot be found by moving from tripartitions to maxcons and then to their induced tripartitions. A strongly stable tripartition requires such a path being a loop of length two (tripartition – maxcon – same tripartition), while the lemma shows a loop of length four.

Incidentally, the sources of the proof have a strongly stable tripartition. Indeed, the only maxcon of $(\{B, C\} || \{A\}, \{D, E\}, \{F, G\})$ is $\{B, C\}$, which induces the partition itself. This means that the tripartition – maxcon – tripartition procedure may fail to find a strongly stable tripartition even if one exists. However, it is also the case that a set of sources may have no strongly stable tripartition.

Theorem 17 *Some sets of sources have no strongly stable tripartitions.*

Proof. The claim is proved using four sources: $\{A, A'\}$, $\{B, B'\}$, $\{C, C'\}$, and $\{D, D'\}$. The eight maxcons are:

$$\begin{aligned} &\{A, A', D, B\} \\ &\{A, A', D', C\} \\ &\{B, B', A, C\} \\ &\{B, B', A', D\} \\ &\{C, C', B, D\} \\ &\{C, C', B', A\} \\ &\{D, D', C, A\} \\ &\{D, D', C', B\} \end{aligned}$$

The first two maxcons are depicted in Figure 11. The second two have a similar graphical representation, and the same for all subsequent pairs.

These maxcons are built in a regular way from the circle of symbols A, B, C, D, A . This can be observed by reading the symbols in vertical, as if they formed four columns: the first column is A, A, B, B, C, C, D, D , the second $A', A', B', B', C', C', D', D'$, the third $D, D', A, A', B, B', C, C'$, the fourth B, C, C, D, D, A, A, B . This means that increasing each symbol (replacing A with B , etc.) does not change these maxcons. For example, if some property is proved on the maxcon $\{A, A', D, B\}$, the same property holds when replacing A with B , A' with B' , B with C , etc.

The first maxcon $\{A, A', D, B\}$ induces $(\{A, A'\} || \{D, D'\}, \{B, B'\} || \{C, C'\})$. The source $\{C, C'\}$ is in the third class because otherwise either C and C' would be consistent with the maxcon $\{A, A', D, B\}$, contradicting its maximal consistency. A similar argument proves that $\{D, D'\}$ and $\{B, B'\}$ are in the second class.

For the same reasons, $\{A, A', D', C\}$ induces $(\{A, A'\} || \{D, D'\}, \{C, C'\} || \{B, B'\})$.

Tripartition $(\{A, A'\} || \{D, D'\}, \{B, B'\} || \{C, C'\})$ have each both maxcons $\{A, A', D, B\}$ and $\{A, A', D', C\}$. The same holds for $(\{A, A'\} || \{D, D'\}, \{C, C'\} || \{B, B'\})$. As a result, they are not strongly stable. By symmetry, the same holds for all three other pairs of maxcons. \square

What to do when no strongly stable partition exist? A weakly stable partition is still coherent to some extent, since at least some of its maxcons induce itself. Yet, some weakly stable partitions may be more coherent than others. The proof of the theorem shows such a case: each partition has a maxset that induces another partition, forming a sort of cycle. Total coherence requires each partition to only induce itself; a milder condition is that a




Figure 11: How the formulae in the proof of Theorem 17 connect. A similar scheme connects B, B', C, D, A', A' , etc.

partition may induce another, but that induces it back. This is not strong stability, but is still more than weak stability. Is like the set of partitions are strongly stable as a group.

Most of the following analysis concerns tripartitions that have a maxcon that induces a possibly different tripartition. Expressing this condition in words every time makes for long and complicated sentences. A formal definition simplifies them.

Definition 12 *The relation \Rightarrow is defined as: $X \Rightarrow Y$ if and only if the tripartition X has a maxcon that induces the partition Y .*

The relation \Rightarrow is defined over all tripartitions, including the ones that are not weakly stable. In particular, X is weakly stable if and only if $X \Rightarrow X$. This implies that \Rightarrow is not reflexive, since $X \Rightarrow X$ does not hold if X is not weakly stable. Restricting to the weakly stable tripartitions makes \Rightarrow reflexive, but neither symmetric nor transitive, as the following theorem proves. The relation is however serial: for every tripartition X it holds $X \Rightarrow Y$ for some possibly different tripartition Y .

Theorem 18 *The relation \Rightarrow is serial but not reflexive, not symmetric and not transitive. When restricted to weakly stable tripartitions it is reflexive but neither symmetric nor transitive.*

Proof. The relation \Rightarrow is serial if $X \Rightarrow Y$ holds for every X and some Y . This is the case because every tripartition X has at least a maxcon, and every maxcon induces a tripartition Y ; therefore, $X \Rightarrow Y$.




Figure 12: These formulae prove that \Rightarrow is not transitive.

Not being reflexive is shown by the tripartition $(\{\{A\}\})$: its only maxcon is $\{A\}$, which induces $(\{\{A\}\})$. Restricting to weakly stable tripartitions makes \Rightarrow reflexive by definition, since weak stability is defined as $X \Rightarrow X$.

The relation \Rightarrow is shown to be not symmetric and not transitive by the formulae in Figure 12. The sources are $\{A\}$, $\{B, C\}$, $\{D, E\}$ and $\{F, G\}$, the plain maxcons $\{A, C, D\}$, $\{A, E, F\}$, $\{A, F, G\}$, $\{B, C\}$ and $\{D, E\}$. The following diagram employs arrows to indicate both that a tripartition has a maxcon and that a maxcon induces a tripartition.

$$\begin{array}{ccc}
 (\{A\}|\{B, C\}, \{D, E\}|\{F, G\}) & \leftrightarrow & \{A, C, D\} \\
 \downarrow & & \uparrow \\
 \{A, E, F\} & \leftrightarrow & (\{A\}|\{F, G\}, \{D, E\}|\{B, C\}) \\
 & & \downarrow \\
 & & \{A, F, G\} \\
 & & \updownarrow \\
 & & (\{A\}, \{F, G\}||\{D, E\}, \{B, C\})
 \end{array}$$

Since A is consistent, the tripartition $(\{A\}|\{B, C\}, \{D, E\}|\{F, G\})$ has only maxcons containing A . The same holds for $(\{A\}|\{F, G\}, \{D, E\}|\{B, C\})$. In particular, both tripartitions have $\{A, C, D\}$ and $\{A, E, F\}$, but only the second has $\{A, F, G\}$. This maxcon induces the tripartition $(\{A\}, \{F, G\}||\{D, E\}, \{B, C\})$. This tripartition has $\{A, F, G\}$ as its only maxcon, showing it to be strongly stable.

In terms of the relation, $(\{A\}|\{F, G\}, \{D, E\}|\{B, C\}) \Rightarrow (\{A\}, \{F, D\}||\{D, E\}, \{B, C\})$ holds but not the other way around, proving that \Rightarrow is not symmetric even on weakly stable tripartitions. Furthermore, the following both holds:

$$\begin{array}{l}
 (\{A\}|\{B, C\}, \{D, E\}|\{F, G\}) \Rightarrow (\{A\}|\{F, G\}, \{D, E\}|\{B, C\}) \\
 (\{A\}|\{F, G\}, \{D, E\}|\{B, C\}) \Rightarrow (\{A\}, \{F, D\}||\{D, E\}, \{B, C\})
 \end{array}$$

Yet, $(\{A\}|\{B, C\}, \{D, E\}|\{F, G\}) \Rightarrow (\{A\}, \{F, D\}|\{D, E\}, \{B, C\})$ does not, proving that the relation is not transitive. \square

The definitions of weak and strong stability can be expressed in terms of the relation \Rightarrow between tripartitions. Indeed, weak stability requires the tripartition to have a maxcon that induces the partition itself ($X \Rightarrow X$) while strong stability also requires that no other tripartition is obtained this way:

weakly stable: $X \Rightarrow X$

strongly stable: $X \Rightarrow Y$ if and only if $X = Y$

Weakly stable tripartitions always exist, but some set of sources do not have any strongly stable tripartition. In terms of the relation \Rightarrow , a tripartition X is weakly but not strongly stable if $X \Rightarrow X$ but also $X \Rightarrow Y$ with $Y \neq X$. In terms of reliability, the reliability information encoded by X produces a maxcon that supports X itself, but also a maxcon that supports another reliability ordering Y . The supported ordering may be X itself or another one Y depending on how formulae are integrated, since a maxcon is a way to integrate formulae according to the tripartition.

In short, X supports itself, but also another ordering Y .

Starting from a reliability ordering X , an alternative ordering Y is obtained. If Y in turn supports X , some kind of stability is still achieved: X admits a different ordering Y , but Y also admits X . Strong stability means that X only justifies itself; a milder form of stability is that X justifies itself, albeit indirectly, no matter how formulae are integrated.

From a different perspective, X and Y are stable as a group. If the mayor could have said “my experts in order of reliability are Akai, Bialetti and Connors, or maybe Akai, Connors and Bialetti”, he would have supported two alternative partitions instead of one. The first ordering leading to a decision that implies the second would not have looked so incoherent, since the second was also supported by the mayor. The same holds for the second ordering leading to a decision supporting the first.

The problem with switching from one partition to more is that more partitions generally have more maxcons; since maxcons are disjointed, the resulting knowledge is weaker. In the example, the mayor really needed a definite solution, a piece of knowledge uniquely telling a plan to take; “we will build a retaining wall *or* remove soil from the head of the slide” is not a solution. A single strongly stable partition always beats a stable group of partitions; yet, Theorem 17 warns that stable partitions may not exist, so resorting to group coherence may be the only available choice.

Definition 13 *A tripartition is mildly stable if it is weakly stable and one of the following alternative conditions hold, where $\overset{*}{\Rightarrow}$ is the transitive closure of \Rightarrow :*

mild-1 stability: *if $X \Rightarrow Y$ then $Y \Rightarrow X$*

mild-2 stability: *if $X \Rightarrow Y$ then $Y \overset{*}{\Rightarrow} X$*

mild-3 stability: *if $X \overset{*}{\Rightarrow} Y$ then $Y \overset{*}{\Rightarrow} X$*

Rather than insisting on X always leading to itself, these conditions allow for X to lead to some other tripartition Y , but only if coming back to X is possible in some way. In particular, mild-1 stability requires that a single step suffices, mild-2 stability allows for a path, and mild-3 stability disallows for a tripartition to be reachable without the original tripartition being reachable from it.

None of the three forms of mild stability imply weak stability, since they only constraint the case where $X \Rightarrow Y$ without saying anything about $X \Rightarrow X$. This is why the definition of mild stability also requires the tripartition to be weakly stable.

Conversely, weak stability does not imply any form of mild stability: $X \Rightarrow X$ only implies that a maxcon M of X induces X itself, but X may also have another maxcon M' that induces a different tripartition Y . In this case, $X \Rightarrow Y$ and $X \xRightarrow{*} Y$, which is the situation the three forms of mild stability constraint.

Mild-2 stability and mild-3 stability may look the same, but they are not; the three forms of stability are all different from each other. In particular, mild-3 stability implies mild-2 stability but not the other way around; for example, if $X \Rightarrow Z$ and $Z \Rightarrow Y$ then mild-3 stability implies $Y \xRightarrow{*} X$ while mild-2 stability only implies $Z \xRightarrow{*} X$ and $Y \xRightarrow{*} Z$. That \Rightarrow is not transitive is proved in Theorem 18.

Theorem 19 *Some sets of sources have no mild-1 stable tripartitions.*

Proof. The sources are $\{A, A', A''\}$, $\{B, B', B''\}$, $\{C, C', C''\}$ and $\{D, D', D''\}$. For each, a group of three maxcons contains it all; for the first source, these are $\{A, A', A'', B, B', C\}$, $\{A, A', A'', C, C', D\}$ and $\{A, A', A'', D, D', B\}$. Another similar group of four maxcons contains all of $\{B, B', B''\}$, etc. These maxcons do not contain each other since: a. each has a whole source and a part of the others; and b. the ones that contain the same whole source (like two that contain all of $\{A, A', A''\}$) have the first two formulae of another source (like $\{B, B'\}$) and one of yet another (like $\{C\}$). Since these sets of symbols do not contain each other, they can be realized by maxcons thanks to Lemma 5.

Since mild-1 stable tripartitions are also weakly stable by definition, only tripartitions induced by one of these maxcons can be mild-1 stable. The first maxcon $\{A, A', A'', B, B', C\}$ induces the following tripartition:

$$(\{A, A', A''\}|\{B, B', B''\}, \{C, C', C''\}|\{D, D', D''\})$$

This tripartition has two maxcons: $\{A, A', A'', B, B', C\}$ and $\{A, A', A'', C, C', D\}$. The first is a maxcon because it induces the tripartition itself; the second is an alternative because it contains C' from the second class, which the first does not. This maxcon induces a different tripartition:

$$(\{A, A', A''\}|\{C, C', C''\}, \{D, D', D''\}|\{B, B', B''\})$$

Again, this tripartition has the maxcon $\{A, A', A'', C, C', D\}$ because it was induced by it, but also $\{A, A', A'', D, D', B\}$ because this maxcon contains D' from the second class, which is not in the previous maxcon. The second maxcon induces the tripartition:

$$(\{A, A', A''\}|\{D, D', D''\}, \{B, B', B''\}|\{C, C', C''\})$$

This tripartition has the maxcon $\{A, A', A'', D, D', B\}$ but also $\{A, A', A'', B, B', C\}$ because the latter has B' while the first does not. This was the maxcon the proof started from. Therefore, a cycle of three tripartition has been shown:

$$\begin{array}{ccc}
 & (\{A, A', A''\}|\{B, B', B''\}, \{C, C', C''\}|\{D, D', D''\}) & \\
 & \swarrow & \uparrow \\
 (\{A, A', A''\}|\{C, C', C''\}, \{D, D', D''\}|\{B, B', B''\}) & & \\
 & \searrow & \\
 & (\{A, A', A''\}|\{D, D', D''\}, \{B, B', B''\}|\{C, C', C''\}) &
 \end{array}$$

The point of the proof is that, apart from $\{A, A', A''\}$, each maxcon has two formulae in a source S and one in another S' ; its induced tripartition has S and S' in the second class. However, another maxcon has two formulae in S' , and is therefore an alternative maxcon of the same tripartition. Its induced tripartition has S in the third class and another source in the second. A mechanism like this allows building cycles of tripartitions and maxcons of arbitrary length.

What holds for the maxcons containing $\{A, A', A''\}$ holds for the maxcons containing $\{B, B', B''\}$, for the maxcons containing $\{C, C', C''\}$ and for the maxcons containing $\{D, D', D''\}$ because these are the same apart from a renaming of the formulae. Therefore, this set of sources has no mild-1 stable tripartition. \square

The mechanism used in the proof can be generalized so that counterexamples and specific scenarios can be easily created. This is in the same line of Lemma 5, which shows that sets of symbols always express maxcons given that no set contain another.

For tripartitions, the construction is not that clean. Ideally, it should be possible to give an arbitrary graph and obtain a set of sources so that their tripartitions are related as in the graph. Instead, the following lemma shows how to obtain a graph only for the weakly stable tripartitions, and only if the graph has a specific form.

Lemma 11 *Given a directed graph of n disconnected subgraphs of n nodes each, a set of sources exists such that the nodes of the graph correspond to its weakly stable tripartitions and the edges to its tripartition–maxcon–tripartition relation \Rightarrow .*

Proof. Let X_j^i be the j -th node of the i -th subgraph. The existence of an edge from X_j^i to X_z^i is denoted by $X_j^i \rightarrow X_z^i$.

The sources are S_0, S_1, \dots, S_n where $S_i = \{F_i, F'_i, F''_i\}$. The first subgraph is obtained by the following maxcons:

$$M_j^0 = \{F_0, F'_0, F''_0, F_j, F'_j\} \cup \{F_z \mid X_j^i \rightarrow X_z^i\} \text{ for every } j \in [1, n]$$

The maxcons for the other subgraphs are described below, for the moment the only important part of the construction is that none of them contains all three formulae F_0 , F'_0 and F''_0 .

The weakly stable tripartitions are those induced by at least a maxcon. The tripartition P_j induced by M_j^0 has $S_0 = \{F_0, F'_0, F''_0\}$ in the first class, $S_j = \{F_j, F'_j, F''_j\}$ in the second together with all $S_z = \{F_z, F'_z, F''_z\}$ such that $X_j^0 \rightarrow X_z^0$. The remaining sources are in its third class.

$$P_j^0 = (S_0 | S_j S_z \dots | \dots)$$

Since S_0 is consistent, only plain maxcons containing all of S_0 can be maxcons of this tripartition P_j^0 . These are: M_j^0 , the maxcons M_z^0 such that $X_j^0 \rightarrow X_z^0$ contain all formulae in the first class, and the other maxcons M_w^0 . These three kinds of maxcons are considered in turn:

- M_j^0 is a plain maxcon and it contains all formulae in the first class and none in the third; by Lemma 8, it is a maxcon;
- let z be such that $X_j^0 \rightarrow X_z^0$; by construction, M_z^0 contains all formulae in the first class and F'_z ; no other maxcon contains all these four formulae; therefore, this one is maximal both in the first class and in the union of the first and the second;
- the other plain maxcons M_z^0 are the only ones containing all formulae in the first class of the tripartition; of the second class they may contain F_j , but neither F'_j (because only M_j does) nor F'_z (only M_z^0 does); the intersection of this maxcon with the first and the second class of the tripartition is therefore strictly contained in M_j^0 .

These three points prove that the maxcons of the tripartition P_j^0 are M_j^0 and all M_z^0 with $X_j^i \rightarrow X_z^i$. By construction, M_j^0 induces P_j^0 while M_z^0 induces P_z^0 . Therefore, every P_j^0 is a weakly stable partition, and $P_j^0 \Rightarrow P_z^0$ if and only if $X_j^0 \rightarrow X_z^0$.

By symmetry, what holds for index $i = 0$ holds for the others. For example, the maxcon M_j^1 induces P_j^1 for every $j \neq i$. What makes all these cases disconnected is that M_j^i contains all three formulae F_i, F'_i, F''_i , while no other M_j^w contains F''_i .

The only caveat to this construction is that for $i = 0$ the three formulae F_0, F'_0, F''_0 are in all maxcons and are the first class of every tripartition. For $i = 1$, the same role is taken by F_1, F'_1, F''_1 . Therefore, the node X_1^1 cannot be represented by P_1^1 and M_1^1 , but by P_0^1 and M_0^1 . In general, for $j \leq i$ the node X_j^i is represented by M_{j-1}^i and P_{j-1}^i . Apart from this change, every node in the graph has a corresponding tripartition, and the edges correspond to the relation among tripartitions. \square

While the graphs that can be realized by using this lemma have the quite specific form of n disconnected subgraphs of n nodes each, they are sufficient for showing some interesting properties. For example, that \Rightarrow is neither symmetric nor transitive even when restricted to weakly stable tripartitions only is proved by the graph in Figure 13.

Contrarily to the previous figures, Figure 13 is a formal proof of the claim. The previous figures were just illustrations of how the involved formulae look like when depicted in the space of models. Figure 13 shows a graph that Lemma 11 proves can be translated into a set of sources, each with actual formulae; their tripartitions correspond to the nodes of this graph and their \Rightarrow relation to its edges. If X, Y and Z are the tripartitions that correspond to the three upper nodes, then $X \Rightarrow Y, Y \Rightarrow Z$ and no other tripartition is in relation with them. In particular, neither $Y \Rightarrow X$ nor $X \Rightarrow Z$ holds; this is a valid proof of the claim thanks to Lemma 11.

Another graph proves that some sources have no mild-1 stable tripartitions. It is shown in Figure 14: it has three disconnected cycles of three nodes each. Again, this is a formal




Figure 13: The relation \Rightarrow is neither symmetric nor transitive even when restricted to weakly stable tripartitions.




Figure 14: Some sources have no mild-1 stable tripartitions.

proof of the claim: by Lemma 11, some set of sources has exactly these weakly stable tripartitions in this \Rightarrow relation. Every node of the graph has a successor that is not connected back to it; therefore, for every tripartition X it holds $X \Rightarrow Y$ but not $Y \Rightarrow X$. This again is a valid proof thank to Lemma 11.

Lemma 11 also allows for proving that mild-1, mild-2 and mild-3 stability are all different from each other. A single graph suffices, since it can be implemented by replicating it a number of times. Since the lemma only concerns weakly stable tripartitions, the loops are omitted.

$$\begin{array}{ccc} X & \Leftarrow & \\ \Downarrow & & Z \Rightarrow W \\ Y & \Rightarrow & \end{array}$$

The relation is defined by: $X \Rightarrow Y$, $Y \Rightarrow X$, $Y \Rightarrow Z$, $Z \Rightarrow X$ and $Z \Rightarrow W$. The corresponding graph has four nodes; therefore, it has to be replicated in four copies to allow for the application of the lemma. Only the first copy is considered, since the others are identical.

The tripartition X is mild-1 stable, since $X \Rightarrow Y$ and $Y \Rightarrow X$. The tripartition Y is not, since $Y \Rightarrow Z$ but not the converse. However, since $Z \Rightarrow X$ and $X \Rightarrow Y$, it follows $Z \overset{*}{\Rightarrow} Y$. Therefore, Y is mild-2 stable but not mild-1 stable. Apart from W , none of the tripartitions is mild-3 stable since, for example, $Y \overset{*}{\Rightarrow} W$ but not the other way around.

This counterexample shows how mildly stable tripartitions can be visualized: mild-1 requires all edges to form cycles of length at most two; mild-2 is the same without constraints on the length; mild-3 can be recast in terms of maximality.

Theorem 20 *The mild-3 stable tripartitions of a set of sources are the maximal tripartitions with respect to the relation $\overset{*}{\Rightarrow}$.*

Proof. By definition, a tripartition X is not maximal if $X \stackrel{*}{\Rightarrow} Y$ for some Y , but not the converse. This is exactly the definition of X being not maximal with respect to $X \stackrel{*}{\Rightarrow} Y$. \square

This theorem implies that every set of sources has some mild-3 stable tripartition, since $\stackrel{*}{\Rightarrow}$ is by construction reflexive and transitive, and the number of tripartitions is finite. The same can be proved for mild-2 stability.

Theorem 21 *Every set of sources has mild-2 stable tripartitions.*

Proof. The relation \Rightarrow is serial: every tripartition X has at least a maxcon, and this maxcon induces a tripartition Y ; this is the definition of a serial relation: for every X there exists a Y such that $X \Rightarrow Y$. Visualizing \Rightarrow as a graph, every node has at least an outgoing edge.

Some tripartitions may not be mild-2 stable. This is the case if $X \Rightarrow Y$ holds but $Y \stackrel{*}{\Rightarrow} X$ does not. In both cases, X and all other tripartitions Z such that $Z \stackrel{*}{\Rightarrow} X$ can be removed from consideration. By assumption, $Y \stackrel{*}{\Rightarrow} Z$ does not hold; such a removal does not change the tripartitions that are reachable from Y .

The same argument can be repeated for Y : either Y is mild-2 stable or some tripartitions can be removed from consideration. Since the number of tripartitions is finite, at some point either a mild-2 stable tripartition is found or only a single tripartition W is left. In such a case, $W \Rightarrow W$ and the tripartition is strongly, and therefore mild-2, stable. \square

6. Multiple Classes

Assessing the sources as reliable, partly reliable and unreliable is qualitative: no specific numbers are involved, such as defining a source reliable if 90% of the information it provides is correct. Also, a division in more than three classes may make sense. Yet, the general framework can still be applied. If the mayor of the town facing the danger of the landslides had twelve expert, each providing from five to ten statements, he might have a better opportunity to defend his choice. For example, the journalist would have a hard time saying that the mayor followed Connors' advice just because one of the seven points in Connors' report was consistent with the final decision. Yet, if the mayor said that Connors was the least reliable of the experts but then building the retaining wall was entirely consistent with all of Connors' statements, the mayor's choice would be hard to justify. In other words, the framework is still applicable in this more general situation: assume a reliability ordering, merge using it and then check if the result is consistent with the initial assumption. Yet, the formalization and the technical results are expected to be different.

The aim of this section is to start the analysis of this more general setting.

The definition of stability as presented in the previous sections is based on a loop of definitions: a. every partition has a set of maxcons; and b. every maxcon induces a partition. A partition is a way of representing an ordering among sources, and a maxcon is a consistent combination of formulae, one among the best combinations according to the ordering. Generalizing, these are respectively an ordering and one of the best ways of integrating information according to the ordering. The mechanism therefore requires:

1. from a reliability assessment of the sources, define how the information they provide can be best integrated; and

2. from the integrated information, assess the reliability of the sources.

The reliability of the sources may be a partition, but also an ordering of the sources, or a numerical evaluation over them. Their integration could be a set of maxcons, but some information combination methods output an ordering over the interpretations.

6.1 Percentage Bounds

The most obvious extension to the method is to divide the sources on the percentage of right formulae they provide. This allows for as many classes as wanted. For example, the first class contains the sources that provided only correct formulae, the second those providing between 90% and 100% of correct formulae, etc.

A problem with *every* method of this kind is that the bounds are arbitrary: why not setting the division at 89% instead of 90%? Why should a source providing 89% of correct formulae be in the third class while one just a little better at 90% in the second? Such seemingly arbitrary bounds are uncommon in formal logic, but often seen in applied sciences. The ubiquitous concept of statistical significance is based on an arbitrary level that is set to 5% in many fields (Craparo, 2007; McDonald, 2014), implicitly assuming that a probability of 95% suffices to believe that the observed data has not been obtained just by chance but because of a real phenomenon. Similarly, a source providing 90% of correct information can be considered generally reliable, and one providing less than 50% of correct information unreliable. In a further generalization these bounds are arbitrary rather than fixed. While this section considers 90% and 50%, any fixed pair of bounds leads to the same technical results.

Let $C(S_i, M)$ be the fraction of formulae in the set S_i that are consistent with the set M . This function is used with a source S_i and a maxcon M . Therefore, $C(S_i, M)$ indicates how reliable the source S_i is assuming that the state of the word is described by the maxcon M .

$$C(S_i, M) = \frac{|\{A \in S_i \mid M \not\models \neg A\}|}{|S_i|}$$

This function allows for evaluating a source given a maxcon M . In particular, for the two bounds 90% and 50%, it produces the partition of the sources ($R|P|U$):

$$\begin{aligned} R &= \{S_i \mid 0.9 \leq C(S_i, M)\} \\ P &= \{S_i \mid 0.5 \leq C(S_i, M) < 0.9\} \\ U &= \{S_i \mid C(S_i, M) < 0.5\} \end{aligned}$$

Some limit cases are of interest. Setting both boundaries at 0% makes P and U empty and R contain all sources. The only induced partition is that containing all sources in its first class; this partition is strongly stable; all plain maxcons are therefore maxcons of a strongly stable partition.

Fixing the boundaries at 100% and 0% makes U empty; R and P respectively contain the reliable and unreliable sources as in bipartitions. The properties of bipartitions carry over to this case, including the guaranteed existence of strongly stable partitions.




Figure 15: The formulae used to prove Theorem 23.

Theorem 22 *There exist some percentage bounds that make all set of sources have strongly stable partitions.*

Proof. When the bounds are 100% and 0%, R contains all sources that only provide correct formulae, P all remaining sources and U is empty. This is the same as a bipartition, and strongly stable partitions always exists for bipartitions as proved by Theorem 8. \square

This theorem proves that strongly stable partitions always exist for a specific pair of bounds. It leaves open the question whether the same holds for all possible pair of bounds.

Even the existence of weakly stable partition is no longer obvious. In the previous section, this was the case because a maxcon inducing a partition is also a maxcon of this partition. The same does not always hold for percentage bounds.

In the settings of percentage bounds, the reliability of a source is given by the number of formulae that are consistent with a maxcon. If a source provides the same formula a second time, and this formula is consistent with a maxcon, the reliability of the source according to that source increases. Contrary to the previous formalization, it makes sense to assume that sources are multisets (bags) rather than sets for formulae.

To simplify the notation, a formula A provided nine times by a source is written A^9 . Such a source can be recast as the set of the all-different formulae $A \vee T_1, \dots, A \vee T_9$, where each T_i is a formula with a single model that does not satisfy any other formula; this may require new variables, forgetting which (Lang, Liberatore, & Marquis, 2003) results in nine copies of A .

Theorem 23 *For some percentage bounds, the partition induced by a maxcon does not have that maxcon.*

Proof. The claim is proved by a counterexample: the bounds are 90% and 50%, the sources are $\{A^9, C\}$ and $\{C^8, D^2\}$, the plain maxcons are $\{A, B\}$, $\{A, C\}$, and $\{C, D\}$. These formulae are shown in Figure 15.

The sources have the following percentage of formulae consistent with the plain maxcon $\{A, C\}$:

- the maxcon is consistent with A but not with B ; the percentage of consistent formulae of $\{A, B\}$ is therefore 90%;
- the maxcon is consistent with C but not with D , 80% of $\{C, D\}$.

The induced partition is therefore $(\{A^9, B\}|\{C^8, D^2\})$. Since $\{A, B\}$ is the only plain maxcon that is consistent with all formulae in the first class of the partition, it is the only maxcon of the partition. Therefore, $\{A, C\}$ induces this partition but is not a maxcon of it. \square

The proof that every set of sources has a weakly stable partition was based on the property that the partition induced by a maxcon has that maxcon. It therefore works no longer for percentage bounds. The claim still holds, but for a different reason, suggested by the above proof: a plain maxcon is not a maxcon of the induced partition because of another maxcon that is “more consistent” with the formulae in the first class of the partition. In turn, this maxcon may not induce a weakly stable partition for the same reason, but at some point a maximum is reached.

Lemma 12 *For every pair of bounds there exists an irreflexive, antisymmetric and transitive ordering \prec among plain maxcons such that, if M is not a maxcon of its induced partition then $M \prec M'$ for some maxcon M' of this partition.*

Proof. The bounds define how a maxcon M induces a partition $(R|P|U)$. In the following proof the percentage bounds are set at 90% and 50% for the sake of explanation, but every other pair would work. The ordering is based on the following function.

$$e(M) = (M \cap (\cup R), M \cap (\cup P)) \text{ where } (R|P|U) \text{ is the partition induced by } M$$

The function $e(M)$ does not have the partition as an argument because the partition is the one induced by M . For example, $e(M')$ is defined in terms of the partition $(R'|P'|U')$ induced by M' .

Let $(X, Y) \prec (Z, W)$ if either $X \subset Z$ or $X = Z$ and $Y \subset W$. The ordering on maxcons is defined by $e(M) \prec e(M')$. This ordering is by construction irreflexive, symmetric and transitive. Remains to prove that if M is not a maxcon of its induced partition then $e(M) \prec e(M')$ for some maxcon M' .

Let M be a plain maxcon that is not a maxcon of its induced partition $(R|P|U)$. Since M is consistent but not a maxcon of $(R|P|U)$, another consistent set of formulae is strictly larger than M within $\cup R$ or within $\cup(R \cup P)$. Such a set can be extended by some formulae to obtain a maxcon M' with the same property: either $M \cap (\cup R) \subset M' \cap (\cup R)$ or $M \cap (\cup(R \cup P)) \subset M' \cap (\cup(R \cup P))$. These conditions can be rewritten as:

1. $M \cap (\cup R) \subset M' \cap (\cup R)$; or
2. $M \cap (\cup R) = M' \cap (\cup R)$ and $M \cap (\cup P) \subset M' \cap (\cup P)$.

Let $(R'|P'|U')$ be the partition induced by M' . In both Case 1 and Case 2, $M \cap (\cup R) \subseteq M' \cap (\cup R)$. By definition of induced partition, if $S_i \in R$ the maxcon M contains at least 90% of the formulae in S_i . Since $M \cap (\cup R) \subseteq M' \cap (\cup R)$, also M' contains at least 90% of S_i , implying that $S_i \in R'$. This holds for every source in R ; therefore, $R \subseteq R'$.

In the first of the two cases above, $M \cap (\cup R) \subset M' \cap (\cup R)$. Since $R \subseteq R'$, it follows that $M \cap (\cup R) \subset M' \cap (\cup R')$. Therefore, $e(M) \prec e(M')$.

In the second case, $M \cap (\cup R) = M' \cap (\cup R)$ and $M \cap (\cup P) \subset M' \cap (\cup P)$. Since $R \subseteq R'$, the equality implies $M \cap (\cup R) \subseteq M' \cap (\cup R')$, which is the first condition for $e(M) \prec e(M')$ to hold.

If a source is in P , then M contains between 50% and 90% of its formulae. Since $M \cap (\cup P) \subset M' \cap (\cup P)$, then all these formulae are also in M' . But M' may also contain other formulae of this source; therefore, this source can be in P' but also in R' . Formally, $P \subseteq P' \cup R'$.

Let F be a formula in $M' \cap (\cup P)$ that is not in $M \cap (\cup P)$, and S_i the source containing it. Since F is in $\cup P$ but not in $M \cap (\cup P)$, it is not in M . From $S_i \in P$, it follows $S_i \in R' \cup P'$. If $S_i \in R'$ then $M' \cap (\cup R')$ contains F while $M \cap (\cup R)$ does not (because $F \notin M$). This alone implies $e(M) \prec e(M')$. Otherwise, F is in $\cup P'$. Therefore, $M' \cap (\cup P')$ contains F while $M \cap (\cup P)$ does not, which is the second condition for $e(M) \prec e(M')$ to hold. \square

This lemma is in a way similar to Theorem 6 for bipartitions, in that it relates a maxcon with a maxcon of its induced partition. Differently than that it establishes such a relation only if the first maxcon is not a maxcon of its induced partition, and only with a single maxcon of the induced partition. It is however sufficient to prove the existence of weakly stable partitions.

Theorem 24 *For every pair of bounds, weakly stable partitions exist for every set of sources.*

Proof. Let M be a plain maxcon and $(R|P|U)$ its induced partition. If this partition has M as a maxcon, it is weakly stable and the claim is proved. Otherwise, by the previous lemma there exists M' such that $e(M) \prec e(M')$. The same argument can be repeated for M' : either its induced partition has M' as a maxcon, or another maxcon M'' exists such that $e(M') \prec e(M'')$. Since this relation is irreflexive, antisymmetric and transitive and the set of maxcons is finite, at some point a maxcon M^* is reached such that $e(M^*) \prec e(M''')$ does not hold for any other maxcon M''' , proving that M^* is a maxcon of its induced partition. \square

6.2 Ordering

A different way to implement the principle of evaluating the reliability of the sources from a possible combination of the formulae is to use orderings for both. In particular, the relative reliability of sources can be encoded as an ordering over sources, and a possible result of merging as a priority order over the propositional interpretations. This scenario is far from being new: in iterated revision, a sequence of formulae ordered by reliability produces an ordering of the propositional interpretations (Rott, 2009). Iterated revision implies that the formulae are acquired in ascending order of reliability, but this assumption can often be lifted: for example, lexicographic revision (Spohn, 1988) starts with the set comprising only the formula of the highest reliability, then adds the others in order if consistent with the current set; this only requires the formulae to be ordered by reliability, regardless of their arrival order. The interpretations satisfying the resulting set are the most likely ones;

a similar set can be constructed by removing them and repeating the procedure, and its interpretations are the second-most-likely ones.

When the reliability of the formulae is unknown, the fixed point mechanism introduced in this article can be applied: an ordering of the sources of the formulae is assumed, merging is performed according to it and the resulting ordering over interpretations is used to assess the reliability of the sources.

The reliability of sources can for example be expressed by a relation \leq and the priority over interpretations by a relation \preceq . When restricted to total orders, these can be recast in terms of functions from sources and interpretations to integers. The reliability of a source S_i is denoted $\text{reliability}(S_i)$, the priority of an interpretation I is denoted $\text{priority}(I)$. They are related by:

reliability \implies **priority**: $\text{priority}(I) = \min\{\text{reliability}(S_i) \mid I \models A \text{ for some } A \in S_i\}$

priority \implies **reliability**: $\text{reliability}(S_i) = \max\{\text{priority}(A) \mid A \in S_i\}$
 where $\text{priority}(A) = \min\{\text{priority}(I) \mid I \in A\}$.

If an interpretation I is satisfied by a formula in a source of reliability 1 and one in a source of reliability 2, then I has priority 1. In general, an interpretation is likely being the actual state of the world if a reliable formula supports it. This explains the minimization.

On the converse, a source providing a formula that is considered untruthful is assessed as unreliable. The evaluation of both sources and interpretations is qualitative: the number of formulae supporting an interpretation does not matter, as well as how many formulae a source provides.

When comparing two reliability orderings expressed in numeric form, normalization is necessary. Indeed, the orderings in which I is evaluated 0 in both while J is respectively evaluated 1 and 2 may be the same if no model is evaluated to 1 in the second. In other words, removing the “empty levels” (Liberatore, 1997; Hunter & Delgrande, 2015) is necessary when comparing two orderings.

An example is shown on the following sources:

$$\begin{aligned} S_1 &= \{x, y\} \\ S_2 &= \{\neg x \wedge y\} \end{aligned}$$

A stable ordering is $\text{reliability}(S_1) = \text{reliability}(S_2) = 0$. All models of either of the three formulae have priority 0, the other model $\{\neg x, \neg y\}$ has priority 1. As a result, all three formulae have reliability 0, and therefore both sources have reliability 0 as well. This is the original reliability ordering, which is therefore stable.

Not all orderings are stable, as shown by $\text{reliability}(S_1) = 0$ and $\text{reliability}(S_2) = 1$. The priority of models is $\text{priority}(\{x, y\}) = \text{priority}(\{x, \neg y\}) = \text{priority}(\{\neg x, y\}) = 0$, which leads to $\text{reliability}(S_2) = 0$.

6.3 Weighted Merge

The rationale of merging by maxcons is to integrate the information the sources provide by gathering as much of it as possible while maintaining consistency; more than one such

a maximal accumulation (a maxcon) may exist: they are all accepted as alternative ways of integrating information (maxcons are disjointed). The principle is to aim at maximal consistent information, the disjunction of maxcons formalizing the alternative between the maximal consistent combinations.

A different approach is to look for a complete description of the state of the world (a model), under the assumption that the sources are more likely right than wrong. A model that is close to the formulae is likely to be that model; a model that is far is unlikely so. More than one closest world may exist; all these alternatives are considered possible (Revesz, 1997; Konieczny & Pérez, 2011; Konieczny et al., 2002, 2004).

The parallel between the two approaches is clear: some formal object representing information is aimed at (a model or a set of formulae); the search is guided by the assumption that sources likely provide valid information (a formula that is either consistent or close to the actual state of the word); more than one best such object may exist, and they all are accepted as alternatives (disjointed, or all included in the final set of models). The similarity allows for carrying over the framework of stable partitions to merging by distance.

A number of different methods for merging by distance exist, but often the distance $d(I, F)$ of a model I to a formula F is the minimal Hamming distance from I and a model of F . Such a model I is good if the distance is low, and bad otherwise. This distance is calculated for all the given formulae F_1, \dots, F_m ; the badness of I is a combination of them, for example a weighted sum:

$$d(I, F_1, \dots, F_m) = w_1 \times d(I, F_1) + \dots + w_m \times d(I, F_m)$$

Models of small badness are close to the reliable formulae, and are therefore to be preferred. For this reason, the result of merging comprises the models of minimal badness. This minimization can be applied to all models or only those satisfying some integrity constraints. The integrity constraints formalize facts known for certainty, so that models conflicting with them can be ruled out.

The role the partitions had in the previous sections is now taken by the weights. They formalize the reliability of the sources: they measure how much each is to be taken seriously when merging. Applying them to the sources results in the maxcons then and in the closest models now.

Such closest model is considered likely to be the actual world. Assuming that it is, it can be used to estimate the reliability of the sources. Indeed, a source that provides formulae that are much different from the actual world is unreliable. A measure of reliability is derived from a closest model, in the opposite way as the closest models are obtained from the weights. This is enough for defining stability.

Let S_1, \dots, S_m be a set of sources, each comprising one or more formulae. Each source S_i has a weight w_i . When merging, all formulae in S_i have weight w_i . This is the same principle of merging with maxcons: a formula is as reliable as the source it comes from.

- Merging sources S_1, \dots, S_m with weights w_1, \dots, w_m produces the models at minimal weighted distance from the formulae in the sources;
- the reliability of a source S_i from a model I is the maximum distance between I and a formula in S_i .

The reason for the asymmetry between the two points is that the reliability of a source is not its most truthful formula but its least: a source is not trustable if it provides some formulae that are very far from the truth, even if it also provides some that are close. Other ways of combining the distances could be used, such as the average.

The principle of operation is to start with a set of weights, determine the models that are the closest to the formulae, and from these compute the reliability of the sources. This should be consistent with the original weights, at least qualitatively: if a source has a weight greater than another, then its resulting distance should be lower than those.

The following example shows that some weights may be weakly stable while others are not. This means that the fixpoint definition allows excluding the second set, and therefore restricting the possible weights.

$$\begin{aligned}
 S_1 &= \{A, B\} \\
 S_2 &= \{C\} \\
 A &= \neg x_1 \wedge \neg x_2 \wedge \neg x_3 \wedge \neg x_4 \wedge \neg x_5 \wedge \neg x_6 \\
 B &= x_1 \wedge x_2 \wedge \neg x_3 \wedge \neg x_4 \wedge \neg x_5 \wedge \neg x_6 \\
 C &= x_1 \wedge x_2 \wedge x_3 \wedge x_4 \wedge x_5 \wedge x_6 \\
 K &= x_1 \wedge x_2 \wedge x_3 \wedge x_4 \wedge \neg x_6 \\
 I &= \{x_1, x_2, x_3, x_4, \neg x_5, \neg x_6\} \\
 J &= \{x_1, x_2, x_3, x_4, x_5, \neg x_6\}
 \end{aligned}$$

The formula K expresses the constraints: only its two models I and J are considered when computing the result of merging. The distance between these models and the formulae are:

$$\begin{aligned}
 d(I, A) &= 4 \\
 d(I, B) &= 2 \\
 d(I, C) &= 2 \\
 d(J, A) &= 5 \\
 d(J, B) &= 3 \\
 d(J, C) &= 1
 \end{aligned}$$

The starting point of the method is to assume some weights for the sources, for example $w_1 = w_2 = 1$. This is an obvious choice: in lack of information indicating which source to trust more, they are trusted the same.

Using these weights, the badness of I is $w_1 \times (d(I, A) + d(I, B)) + w_2 \times d(I, C) = 1 \times (4 + 2) + 1 \times 2 = 8$. The badness of J is $1 \times (5 + 3) + 1 \times 1 = 9$. The result of merging is therefore I only. This allows for an easy evaluation of the sources: 4 for S_1 and 3 for S_2 . These reliability measures are different, contradicting $w_1 = w_2$. The original weights are not stable.

A different outcome is obtained from $w_1 = 1$ and $w_2 = 2$. The badness of I is $1 \times (4 + 2) + 2 \times 2 = 10$ and that of J is $1 \times (5 + 3) + 2 \times 1 = 8 + 2 = 10$. Both models are in the result of merging, so both have to be considered when checking stability. The distance between I and S_1 is 4, between I and S_2 is 2. These compare the opposite of w_1 and w_2 , as required (large weight means high reliability, which in turns should produce a lower distance). This proves that these weights are weakly stable. A similar calculation for J shows that the same happens for it, so the weights are strongly stable.

7. Discussion and Comparison with Related Work

What to do with stable partitions?

Regardless of bipartitions or tripartitions or multiple-classes partitions, regardless of weak or strong or mild stability, and even regardless of maxcon-based integration or weighted distance merging or whatever other form of merging is used, regardless of all this some reliability orderings are discarded and some others are selected. This selection guides the fusion of the given information. This was the original motivation, but stable partitions provide more than that.

Explicitly stating the input and output of the whole process clarifies this point:

Input: set of sources. Output: their merge. Given a set of sources $S = \{S_1, \dots, S_n\}$, their stable partitions are computed (stability could be weak, strong or mild; partitions could be of two, three or more classes). Each partition P_i represents a reliability ordering that produces a consistent combination of formulae M_i that is coherent with this ordering. Disjoining all these M_i reflects the alternative among all these orderings. The whole process generates a single formula $\bigvee M_i$ out of the sources $S = \{S_1, \dots, S_n\}$. It can be seen as a merging operator, so that for example one can check whether it satisfies the postulated for merging or arbitration (Revesz, 1997; Liberatore & Schaerf, 1998; Konieczny & Pérez, 2011).

Input: set of sources; partial reliability ordering. Output: their merge. The relative reliability of the sources may not be completely known; for example, a source S_1 is more reliable than S_2 and S_3 , but which of S_2 and S_3 is more reliable is not known. Such a scenario is dealt with by simply removing from consideration all stable partitions that conflict with the known reliability information, for example $(S_2|S_1S_3)$.

Input: set of sources. Output: their merge; set of stable partitions. The process of merging as done in the two scenarios above is a black box: put in the sources, get the result. The stable partitions provide an explanation for the result. They can also be used to perform further integrations or even take actions that are outside the merging process itself, as detailed below.

These three input-output configurations are discussed in order. Of course, the second and the third can be combined into a fourth where the input includes a partial specification and the output the set of stable partitions.

7.1 A Merging Operator

The entire process of selecting the stable partitions and disjoining all maxcons obtained from them is a form of merging: the input is a set of sources, the output a way of integrating them while retaining consistency.

The degenerate case of singleton sources and bipartitions recovers the ancient method of inference from maximal consistent subsets (Rescher, 1964; Brewka, 1989; Nebel, 1992, 1998), as proved by Theorem 2. Just like revision can be seen as union followed by consistency recover (Hansson, 1999), merging can be done by collecting all knowledge and then restoring consistency. Singleton sources and bipartitions give exactly this kind of knowledge integration. But the general framework allows more than that. Allowing sources to provide more than one formula already makes a difference: some partitions are excluded, and the final result is built from a smaller set of maxcons. Since maxcons are disjoined, the result is more specific and therefore carries more information.

Seeing the whole process of selecting the stable partitions and their maxcons as a merging operator makes it comparable to existing ones. For example, one may play with the many postulates that have been introduced for belief merging (Revesz, 1997; Liberatore & Schaerf, 1998; Konieczny & Pérez, 2011), and see that the operator obtained from strongly stable bipartitions with empty integrity constraints ($\mu = \text{true}$) obeys postulates 1, 2, 4 and 5 but not 3 and 6 among the ones proposed by Konieczny and Pérez (2011). But then, if one wanted a given set of postulates to be satisfied, they could just employ a method for combining formula that already satisfy them. While such a comparison with existing merging operators may be of some interest, the point of the new framework is not to introduce yet another specific method for combining sets of formulae, but rather to include in the process a check of the relative reliability of the sources. The following two sections make this point explicit.

7.2 Partially Known Reliability

Stable partitions allow for a partial order to be included as an input in the process. If a certain source is known to be reliable, all reliability orderings must have it in their first class. Looking at it from the reverse direction, only stable partitions that have that formula in their first class are considered. Other constraints may be included in the same way; for example, two sources are known to be equally reliable (only stable partitions including them in the same classes are considered) or one being strictly more reliable than another (the class of the first precedes that of the second).

The constraints may be so strict that a strongly stable partition satisfying them does not exist. In such cases, weakly stable partitions may be considered, and if these do not exist the reliability constraints may be rejected as inconsistent with the information the sources provide.

The full meet revision (Alchourrón, Gärdenfors, & Makinson, 1985) is obtained as a particular case. Revising K by A is obtained by constraining A to be in the first class of the stable partitions of the sources $\{K\}$ and $\{A\}$. While full meet revision only produces either $K \wedge A$ or A , the stable partition also explains the choice: $(AK|)$ in the first case and $(A|K)$ in the second. The Ginsberg-Fagin-Ullman-Vardi (Ginsberg, 1986; Fagin, Ullman, & Vardi, 1983) revision can also be obtained in the same way: $\{K_1, \dots, K_n\}$ revised by A

is the disjunction of the maxcons of the stable partitions where $\{A\}$ is in the first class and the other sources are $\{K_1\}, \dots, \{K_n\}$. These encodings are quite obvious and only prove that restricting the stable partitions based on given constraints makes sense.

In the degenerate case where the partitions are constrained to be a specific one the classical method of inference from preferred subset is recovered. The present framework can therefore be seen as an extension of that system, which applies (in full or restricted) to hypothetical reasoning (Rescher, 1964), default logic (Brewka, 1989), iterated revision (Spohn, 1988), and database update (Fagin et al., 1983).

7.3 After Merging Ends

Partitions that are not stable can be used for merging, but the result they produce contradicts them: a source they rank unreliable is instead reliable as far as can be seen from the result of integration, for example. Selecting only the stable partitions not only reduces the number of possible ways of integrating information, but also explains the selection. Rather than just “the result of merging S_1 , S_2 and S_3 is $x \wedge (\neg y \vee z)$ ”, a system using this method could answer “the result of merging S_1 , S_2 and S_3 is $x \wedge (\neg y \vee z)$ because either S_1 is more reliable than S_2 and S_3 , or they are all equally reliable”.

Outputting the selected reliability ordering does not only justify the result. It is information that can also be used outside this particular merge, or even outside of the merging system altogether.

Assuming for example that $(S_1|S_2S_3)$ and $(S_1S_2S_3|)$ are the only stable partitions, a further merging of formulae from the same sources can be preliminarily performed employing these two partitions, saving the time needed to select the new stable partitions. Only once in a while a new re-assessment of the relative reliability of the sources is done, using the last partitions known to be stable in the meantime. Alternatively, one may consider an initial interval of training when stable partitions are selected, and never changed afterwards.

This strategy assumes that after a given number of mergings the stable partitions reflect the real reliability of the sources. As a result, one may also refine the set of partitions, removing the ones that turn out not to be stable but never adding new ones. At the end of this process a single stable partition tells how reliable the sources are.

Even when several partitions are stable, or even when all plain maxcons are maxcons of stable partitions, having a selection of partitions may be relevant outside of the merging process. For example, if a source that represents information coming from a sensor is in the last class of all stable partitions, this indicates that the sensor is faulty. Some action (e.g., replacement, or just exclusion of data coming from it) may then be taken.

8. Directions

The principle studied in this article is that some reliability orderings can be excluded on the ground that using them for merging leads to different orderings. This principle is implemented by assuming sources to generate a single or a set of formulae, that the reliability ordering is represented by a partition of the sources or by weights, and that formulae are integrated by maxcons or by a weighted sum of distances.

Apart from how the reliability ordering and the merging process are realized, the principle itself can be extended in many ways.

8.1 Sure Information

This is possibly the direction of most interest: some sources of unknown reliability provide formulae to be merged, but some information is also known for certain. This is a common scenario in belief merge, encoded by integrity constraints (Konieczny & Pino Pérez, 1999) as follows: only the models that satisfy these constraints are considered when selecting the ones that make the result of merging.

The fixpoint definition allows for a smooth encoding of formulae that are considered certain. They are collected in a single source, and only reliability orderings having that source in the first class are considered. Since the first class of an induced partition contains only formulae that are consistent with the result of merging, this makes such result to be consistent with all these formulae.

Formalizing the certain information as a source has the advantage of a uniform treatment of it along with the other formulae. There is no prior reason why another source could not be equally certain; if so, it is in the first class of the reliability partition, together with the source encoding the sure formulae. In fact, the integrity constraints come from somewhere, the only difference with the other formulae being that their sources can never be considered unreliable, not even partially so.

The reason why this variant may be the case of most interest stems from having certain information to rule out some reliability orderings. The scheme with no integrity constraint uses only the regular sources to reckon their reliability ordering. The integrity constraints ground this evaluation on information that is known for certain.

8.2 Partial Reliability Ordering

Limiting to reliability orderings with a given source in the first class is only a particular case of partial information over the reliability ordering. If S_0 is such a source, $S_0 \leq S_i$ holds for every other source S_i .

More generally, an arbitrary partial knowledge of the ordering can be assumed. For example, it may be that only $S_1 < S_2$ and $S_2 \leq S_3$ are known. This rules out the partition $(S_2|S_1S_3)$ but neither $(S_1|S_2S_3)$ nor $(S_1|S_2|S_3)$. Since bipartitions and tripartitions set a strong division between classes (all true, some true, none true), constraints of type \leq are better suited for them, as otherwise an apparently harmless statement like “ S_1 is strictly more reliable than S_2 ” would turn “ S_1 is only partly reliable” into the strong implication that S_2 is totally unreliable. Such strict constraints like $S_1 < S_2$ are better left to multiple-classes partitions, where they only state that S_2 provides fewer correct formulae than S_1 .

8.3 Order Compatibility

The mayor chooses an order of reliability of the experts: Akai is the more reliable, Bialetti is less so, Connors even less. A rational way of incorporating their reports according to this order is to start with everything Akai wrote, integrating it with whatever Bialetti wrote that would not lead to contradiction; Connors is less reliable than Bialetti, so the information he provides is added only at the end, if not inconsistent with the integration of Akai’s and Bialetti’s reports.

The result of this process is a unified report, which is taken to be the correct information about the landslide issue. In the other way around, knowledge of the actual state of matters allows evaluating the experts: the one that provides only correct information is more reliable than the one who only provides some, in turn more reliable than the one who provided none.

Such a process allows for a nice fixpoint definition: three classes of assumed reliability of the expert, three classes of estimated reliability of the experts. Equality is coherence, inequality is incoherence.

Yet, the assumed and the evaluated reliability orderings are of different kinds. The third class of the assumed ordering is for experts that are less reliable than all others; the third class of the evaluated ordering is for experts that did not provide any correct information. Yet, inequality is still incoherence: if Bialetti is assumed more reliable than Connors, how could Connors give some correct information while Bialetti do not?

This shows that while the three classes–three classes framework is mathematically nice, nothing forbids partitions of differing number of classes. If the mayor had five experts, he could have ordered them into five classes; the journalist would have a point if the expert in class three did not say anything correct, while that in class five did. This comparison is between a partition of five class (assumed reliability order) and one of three (all correct, some correct, nothing correct), but is not equality. It is a form of compatibility: the evaluated ordering should not contradict the assumed.

8.4 Untrustable Sources

The mayor claimed that Connors is the least reliable of the three experts. But least reliable is not the same as untrustable. Least reliable means that the information Connors provides gives precedence to that from Bialetti and Akai. Untrustable means that Connors' report should go straight to the trashcan. In this particular scenario, the mayor would have had a hard time justifying why he totally excluded the report from a reputed expert; instead, making a choice among conflicting opinions is exactly what a politician is supposed to do.

In other scenarios, untrustable sources are common: information coming from the Internet may be unreliable or totally wrong; in everyday discussions, people may intentionally provide false information to obtain the result they want; disinformation is an established mean to discredit political opposition. In belief revision, such scenarios lead to the concept of strategy-proofness (Everaere, Konieczny, & Marquis, 2007).

Two more classes can be defined: unreliable sources and untrustworthy sources. The first provide no useful information; the second provide only false information. Rationally, merging should exclude the first and consider the negation of the second. When evaluating the result of merging, an unreliable source is expected to provide less correct information than a reliable one. However, the policy for dealing with untrustworthy sources is not that easy to define, since they may provide correct information to disguise themselves.

8.5 Flexible Bounds

Dividing the sources on 100% and 0% of correct formulae, as in bipartition and tripartitions, appeals to a qualitative notion of reliability and unreliability. Every other bound (like 90% or 50%) is difficult to motivate over others (like 89% and 51%).

A possible solution is to assume that a partition like $(S_1|S_2S_3|S_4)$ only means that there exists a reason to assume that S_1 is more reliable than both S_2 and S_3 , which are both more reliable than S_4 , without requiring two fixed borders between them. As an example, this partition is acceptable if the four sources provide 80%, 72%, 69% and 30% correct formulae. The bound can be fixed at 75% and 50% to support this division, but also at 79% and 60%, or at 79% and 40%.

This mechanism lifts the decision of an exact pair of percentages of correct formulae that define a source reliable or partly reliable. A partition like $(S_1|S_2S_3|S_4)$ is then stable if such a pair makes the partition itself induced from the result of merging.

8.6 Technical Results

Some other possible expansions of this work regard the study of the existing definitions, rather than the extensions of them. Lemma 11 allows for the creation of a scenario with given partitions and maxcons, but only if these satisfy certain conditions. It would be interesting to find out whether these constraints could be lifted.

The two algorithms presented in this article cover the case of bipartitions only. Extending them to tripartitions or even to the case of arbitrary partitions would make them applicable to more general situations. Also, the complexity of inferring from the maxcons of the stable partitions, or checking for the existence of strongly stable tripartitions, is also an open problem.

9. Recap

This section contains a short summary of the technical results in this article. The main setting is: some sources provide possibly conflicting information; in lack of reliability information, every order of reliability can potentially be assumed; however, merging according to it may produce a result that supports a different reliability ordering.

Sources are formalized as sets of formulae; for example, $S_1 = \{F_5, F_3, F_1\}$ is a source providing three formulae. The relative reliability of sources is described by an ordered partition (Stanley, 1997); for example, $(S_1|S_2S_3)$ means that S_1 is more reliable than S_2 and S_3 , while S_2 and S_3 are equally reliable. These definitions are made formal in Section 2, together with the kind of merging used in most of this article. In particular, given an ordered partition of the sources, a *maxcon* is a possible way of integrating their formulae: is a maximal consistent subset of a set of formulae.

The framework introduced in this article is first put at work on a limit case in Section 3: every source contains exactly one formula, and sources are partitioned in two reliability classes only: reliable or not. Such a reliability ordering is *stable* if using it for merging and then using the result for assessing the reliability of the sources produces the original partition. The following results are obtained:

- not all partitions are stable;
- every set of sources has at least one stable partition;
- the maxcons of the stable partitions are the maxcons obtained from an all-equal ordering of the sources (plain maxcons).

The last result can be seen as confirmatory: a formula should be more or less taken into consideration depending on how truthful the others provided by the same source are; if every source provides a single formula these formulae should be equally believed.

In Section 4, reliability orderings are still two-class partitions, but sources may provide more than one formula each. A partition may have multiple maxcons; all, some or none of them may induce the partition itself, which is respectively strongly stable, weakly stable or unstable. The results obtained for this setting are:

- some weakly stable partitions are not strongly stable;
- the maxcons of the weakly stable partitions are the plain maxcons of the set comprising all formulae provided by all sources;
- a weakly stable partition may lead via a maxcon to another partition; that partition has a strictly larger first class;
- the maximal weakly stable partitions w.r.t. containment of their first class are the strongly stable partitions;
- the maxcons of the strongly stable partitions are equivalent to the maxcons of a certain partition.

Partitions are related by maxcons: the maxcon of a partition induces a possibly different partition. This relation allows for an equivalent definition of strong stability.

Section 5 still maintains sources providing multiple formulae, but partitions them in three classes rather than two. These divisions are called tripartitions to distinguish them from the bipartitions of the previous sections; this is necessary because some results relate bipartitions and tripartitions. The following results are obtained:

- a plain maxcon is also a maxcon of its induced tripartition;
- weakly stable bipartitions correspond to weakly stable tripartitions;
- this is not the case for strong stability;
- the property of containment of the first class of related bipartitions does not extend to tripartitions;
- rather, cycles of related tripartitions are possible;
- some sets of sources do not have strongly stable tripartitions;
- a third kind of stability can be defined as maximality according to the relation among partitions.

Maximal tripartitions may not be strongly stable, as opposed to the case of bipartitions. Defining mild stability as maximality according to the relation between tripartitions, every source has at least a mildly stable partition, even if it has no strongly stable partition.

Section 6 sticks to the principle of assuming a reliability ordering and then matching it with the result of merging, but extends the analysis in several directions:

- divide the sources according to the percentage of correct formulae they provide;
- use an arbitrary order of the sources for expressing reliability;
- merge by weights.

Section 8 explores future directions. One looks particularly applicable in practice: one or more sources are known for certain to be completely reliable. More generally, the reliability ordering may be partially known. Other cases considered in this section are that of unreliable sources considered as intentionally providing false information and non-fixed percentage bounds.

10. Conclusions

The principle of excluding some reliability orderings when the information integration they produce conflicts with them can be implemented in many ways. In the simplest case, sources are classified in two reliability classes, or three. This seemingly small difference has a big impact on the technical results: strongly stable bipartitions always exist, strongly stable tripartitions do not. This motivates the definition of an ordering among weakly stable partitions in order to obtain a form of stability that is more grounded than weak stability but still has guaranteed existence.

Switching to classes based on percentages makes another important property fail: the partition induced by a maxcon may not have that maxcon. The existence of weakly stable partitions for arbitrary pairs of bounds, which was a simple consequence of this fact, had to be proved by a different argument.

All of this indicates that the technical consequences are strongly dependent on the details of the definition employed. While the general principle is the same (fix a reliability ordering, integrate information using it, check whether the result agrees with the reliability ordering), the properties of the framework depend on how it is implemented.

However, some general facts appear to hold regardless:

1. weakly stable partitions always exist; this is not obvious from the definition, rather the opposite: a definition of partitions and merge that does have this property is unusable since weak stability is the most basic requirement for merging; if some sources do not have such a partition, they cannot be merged; it is hard to imagine a scenario where information cannot be integrated, not even by a void result;
2. since maxcons induce partitions and partitions have maxcons, an ordering of the partitions (or maxcons, or both) always exists; such an ordering can be used to define a form of maximality of partitions, or at least to exclude partitions that are strictly dominated by others; the rationale is that if a reliability ordering lead to an information integration that implies a different reliability ordering, the latter is more grounded than the former;
3. as in the general problem of belief revision, the aim is at the same time to obtain some merging result but also to restrict the possible alternatives as much as possible, since this leads to a resulting formula that carries as much information as possible; as a

result, whichever pair of definitions are used for the combination of formulae and the induced reliability ordering, some stable partitions should always exist, but as few as possible; this is why definitions that lead to cases without weakly stable partitions are ruled out; at the same time, an inflation of weakly stable partitions (like when they lead to all plain maxcons) motivates the introduction of strong stability or at least the restriction to maximal weakly stable partitions.

All these results are relative to how this article implements the general mechanism of assume-and-verify for reliability orderings. Section 6 and Section 8 hint that other realizations of it are possible. This article presents both an idea and a possible implementation, but other implementations are possible. The idea is to assume a reliability ordering and then check it against the result it produces; the implementation in this article is to the case of integrating by maxcons, with consistency used for assessing reliability. Other implementations are certainly possible. It would be of not surprise if the reader already at page ten had some ideas about how to differently instantiate the general framework to better fit a particular domain or example scenario.

Few articles in the literature are related to the research presented in this article. This is because most work in belief revision and merging is about going from a set of formulae to their aggregation. A handful of articles consider the reverse problem: Haret, Mailly, and Woltran (2016) divide a formula into simpler formulae that give the original one when merged; Booth and Nittka (2008) and Liberatore (2015) use the history of previous revision to obtain reliability information; Liberatore (2016) uses merging samples to the same aim. All this research shares the reversal of the traditional role of the reliability information and the result of merge, where the reliability is input and the merging result is output. Apart from the first cited article, which aims at representing knowledge in some simple form, the others do this to obtain reliability information, which can be then used in the following merging. The present article does the same, but does not assume any prior knowledge. Rather, the role of history or examples is taken by the merging result itself, which allows to evaluate the sensibility of the reliability ordering initially assumed.

Acknowledgements

The author thanks the reviewers for their valuable suggestions.

References

- Adamic, L., Zhang, J., Bakshy, E., & Ackerman, M. (2008). Knowledge sharing and Yahoo Answers: everyone knows something. In *Proceedings of the Seventeenth International Conference on World Wide Web (WWW'08)*, pp. 665–674.
- Alchourrón, C. E., Gärdenfors, P., & Makinson, D. (1985). On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50, 510–530.
- Baral, C., Kraus, S., Minker, J., & Subrahmanian, V. (1992). Combining knowledge bases consisting of first-order theories. *Computational Intelligence*, 8(1), 45–71.

- Barua, A., Thomas, S., & Hassan, A. (2012). What are developers talking about? An analysis of topics and trends in Stack Overflow. *Empirical Software Engineering*, 19(3), 619–654.
- Benferhat, S., & Yahi, S. (2009). Complexity and cautiousness results for reasoning from partially preordered belief bases. In *Proceedings of the tenth european conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU'09)*, pp. 817–828.
- Bonatti, P., Calimeri, F., Leone, N., & Ricca, F. (2010). Answer set programming. In *A 25-Year Perspective on Logic Programming*, chap. 5. Springer.
- Booth, R. (2006). Social contraction and belief negotiation. *Information Fusion*, 7(1), 19–34.
- Booth, R., & Hunter, A. (2018). Trust as a precursor to belief revision. *Journal of Artificial Intelligence Research*, 61, 699–722.
- Booth, R., & Nittka, A. (2008). Reconstructing an agent’s epistemic state from observations about its beliefs and non-beliefs. *Journal of Logic and Computation*, 18, 755–782.
- Brewka, G. (1989). Preferred subtheories: an extended logical framework for default reasoning. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence (IJCAI'89)*, pp. 1043–1048.
- Cevolani, G. (2014). Truth approximation, belief merging, and peer disagreement. *Synthese*, 191(11), 2383–2401.
- Craparo, R. (2007). Significance level. In *Encyclopedia of measurement and statistics*, Vol. 3, pp. 889–891. Sage.
- D’Alfonso, S. (2016). Belief merging with the aim of truthlikeness. *Synthese*, 193(7), 2013–2034.
- Delgrande, J., Dubois, D., & Lang, J. (2006). Iterated revision as prioritized merging. In *Proceedings of the Tenth International Conference on Principles of Knowledge Representation and Reasoning, (KR-2006)*, pp. 210–220.
- Everaere, P., Konieczny, S., & Marquis, P. (2007). The strategy-proofness landscape of merging. *Journal of Artificial Intelligence Research*, 28, 49–105.
- Everaere, P., Konieczny, S., & Marquis, P. (2015). Belief merging versus judgment aggregation. In *Proceedings of the Fourteenth International Conference on Autonomous Agents and Multiagent Systems (AAMAS'15)*, pp. 999–1007.
- Fagin, R., Ullman, J. D., & Vardi, M. Y. (1983). On the semantics of updates in databases. In *Proceedings of the Second ACM SIGACT SIGMOD Symposium on Principles of Database Systems (PODS'83)*, pp. 352–365.
- Gabbay, D., & Schlechta, K. (2016). *Reiter Defaults and Autoepistemic Logic*, chap. 4. Springer.
- Ginsberg, M. L. (1986). Counterfactuals. *Artificial Intelligence*, 30, 35–79.
- Hansson, S. (1999). A survey of non-prioritized belief revision. *Erkenntnis*, 50(2), 413–427.

- Haret, A., Maily, J.-G., & Woltran, S. (2016). Distributing knowledge into simple bases. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence (IJCAI 2016)*, pp. 1109–1115.
- Herzig, A., Pozos-Parra, P., & Schwarzentruher, F. (2014). Belief merging in dynamic logic of propositional assignments. In *Proceedings of the eighth International Symposium on the Foundations of Information and Knowledge Systems (FoIKS 2014)*, pp. 381–398. Springer.
- Hunter, A., & Delgrande, J. (2015). Change with uncertain action histories. *Journal of Artificial Intelligence Research*, 53, 779–824.
- Katz, Y., & Golbeck, J. (2006). Social network-based trust in prioritized default logic. In *Proceedings of the Twenty-First National Conference on Artificial Intelligence (AAAI 2006)*, pp. 1345–1350.
- Konieczny, S. (2004). Belief base merging as a game. *Journal of Applied Non-Classical Logics*, 14, 275–294.
- Konieczny, S., Lang, J., & Marquis, P. (2002). Distance-based merging: a general framework and some complexity results. In *Proceedings of the Eighth International Conference on Principles of Knowledge Representation and Reasoning (KR 2002)*, pp. 97–108.
- Konieczny, S., Lang, J., & Marquis, P. (2004). DA^2 merging operators. *Artificial Intelligence*, 157(1-2), 49–79.
- Konieczny, S., & Pérez, R. (2011). Logic based merging. *Journal of Philosophical Logic*, 40(2), 239–270.
- Konieczny, S., & Pino Pérez, R. (1999). Merging with integrity constraints. In *Proceedings of the Fifth European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty, (ECSQARU'99)*, pp. 233–244.
- Lang, J., Liberatore, P., & Marquis, P. (2003). Propositional independence - formula-variable independence and forgetting. *Journal of Artificial Intelligence Research*, 18, 391–443.
- Liberatore, P. (1997). The complexity of belief update. In *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence (IJCAI'97)*, pp. 68–73.
- Liberatore, P. (2015). Revision by history. *Journal of Artificial Intelligence Research*, 52, 287–329.
- Liberatore, P. (2016). Belief merging by examples. *ACM Transactions on Computational Logic*, 17(2).
- Liberatore, P., & Schaerf, M. (1998). Arbitration (or how to merge knowledge bases). *IEEE Transactions on Knowledge and Data Engineering*, 10(1), 76–90.
- McDonald, J. (2014). *Handbook of Biological Statistics* (third edition). Sparky House Publishing.
- Naumann, F., Bilke, A., Bleiholder, J., & Weis, M. (2006). Data fusion in three steps: Resolving schema, tuple, and value inconsistencies. *IEEE Data Eng. Bull.*, 29(2), 21–31.

- Nebel, B. (1992). *Syntax-Based Approaches to Belief Revision*, pp. 52–88. Cambridge University Press.
- Nebel, B. (1998). How hard is it to revise a belief base?. In Dubois, D., & Prade, H. (Eds.), *Belief Change - Handbook of Defeasible Reasoning and Uncertainty Management Systems, Vol. 3*. Kluwer Academic.
- Posnett, D., Warburg, E., Devanbu, P., & Filkov, V. (2012). Mining Stack Exchange: expertise is evident from initial contributions. In *Proceeding of the 2012 International Conference on Social Informatics (SocialInformatics)*, pp. 199–204.
- Reiter, R. (1980). A logic for default reasoning. *Artificial Intelligence*, 13, 81–132.
- Rescher, N. (1964). *Hypothetical reasoning*. North-Holland.
- Revesz, P. (1997). On the semantics of arbitration. *International Journal of Algebra and Computation*, 7, 133–160.
- Rott, H. (1993). Belief contraction in the context for the general theory of rational choice. *Journal of Symbolic Logic*, 58(4), 1426–1450.
- Rott, H. (2009). Shifting priorities: Simple representations for twenty-seven iterated theory change operators. In Makinson, D., Malinowski, J., & Wansing, H. (Eds.), *Towards Mathematical Philosophy*, Vol. 28 of *Trends in Logic*, pp. 269–296. Springer Netherlands.
- Seattle DPD (2000). Landslide study. Geotechnical report.
- Shackel, N. (2007). Bertrand’s paradox and the principle of indifference. *Philosophy of Science*, 74(2), 150–175.
- Spohn, W. (1988). Ordinal conditional functions: A dynamic theory of epistemic states. In *Causation in Decision, Belief Change, and Statistics*, pp. 105–134. Kluwer Academic.
- Stanley, R. (1997). *Enumerative Combinatorics*. Cambridge University Press.
- Viégas, F., Wattenberg, M., Kriss, J., & van Ham, F. (2007). Talk before you type: coordination in Wikipedia. In *Proceeding of the Fortieth Hawaii International Conference on Systems Science (HICSS-40)*, p. 78.
- Wikipedia (2006). Porchesia. <http://wikimedia.7.x6.nabble.com/Porchesia-td962303.html>.
- Wiley, D., & Gurrell, S. (2009). A decade of development. *Open Learning: The Journal of Open, Distance and e-Learning*, 24, 11–21.