Available online at www.sciencedirect.com

SCIENCE DIRECT®

Information Fusion 7 (2006) 19–34

INFORMATION FUSION

www.elsevier.com/locate/inffus

Social contraction and belief negotiation ☆

Richard Booth *

Department of Computer Science, University of Leipzig, Augustusplatz 10/11, 04109 Leipzig, Germany

Received 7 October 2003; revised 15 December 2004; accepted 9 January 2005
Available online 12 February 2005

Abstract

An intelligent agent may receive information about its environment from several different sources. How should the agent merge these items of information into a single, consistent piece? Taking our lead from the contraction + expansion approach to belief revision, we envisage a two-stage approach to this problem. The first stage consists of weakening the individual pieces of information into a form in which they can be consistently added together. The second, trivial, stage then consists of simply adding together the information thus obtained. This paper is devoted mainly to the first stage of this process, which we call *social contraction*. We consider both a postulational and a procedural approach to social contraction. The latter builds on the author's framework of *belief negotiation models*. With the help of Spohn-type rankings we provide two possible instantiations of this extended framework. This leads to two interesting concrete families of social contraction functions.

Keywords: Merging; Belief contraction; Spohn rankings; Multi-agent systems

1. Introduction and preliminaries

An intelligent agent may receive information about its environment from several different sources. How should the agent *merge* these pieces of information into a single, consistent piece? This question has recently received various treatments (see e.g. [5,7,12,13,15–17,20]). The simplest thing to do would be to just take the given pieces of information and *conjoin* them. While this strategy would be fine if the pieces of information are jointly consistent, it could well be that some of the pieces stand in contradiction, in which case the strategy breaks down. In this paper we envisage a two-stage approach to the problem: (i) the individual, raw pieces of information are manipulated (more precisely, *weakened*) into a form in which they *become* jointly consistent, and *then* (ii) the pieces thus obtained are conjoined. Stage (ii) is trivial. Stage (i) is not, and so forms the main topic of this paper.

A precedent for this two-stage approach can be found in the literature on the closely-related area of *belief revision* [1,8,11]. Belief revision may essentially be thought of as "binary merging". It addresses the problem of how to merge one item of information, usually taken to represent the current beliefs of some agent, with another item, representing some new piece of information which the agent acquires. The idea, which dates back to [14] and is given succinct expression by the *Levi Identity* [8], is that this operation of revision is decomposed into two sub-operations: (i) *contraction*: the current information is weakened so that it becomes consistent with the new information, then (ii) *expansion*: the new information is simply added to the result. Note that, in (i), only the current information is weakened, not the new. This reflects the traditional assumption that the new information is always *completely reliable*. What we seek in this paper is a generalised version of the contraction operation. One in which several items of information may all be weakened *simultaneously* so that they become

☆ Expanded version of a paper presented at the Eighth Conference on Principles of Knowledge Representation and Reasoning (KR 2002).
* Tel.: +49 341 9732282; fax: +49 341 9732299.
E-mail address: booth@informatik.uni-leipzig.de

1566-2535/$ - see front matter © 2005 Elsevier B.V. All rights reserved.
doi:10.1016/j.inffus.2005.01.002

consistent *with one another*. For this reason we call the operations we are interested in *social contraction* functions (SC functions for short).

We shall examine social contraction from two viewpoints: a *postulational* one and a more *procedural* one. For the latter we build on the framework of *belief negotiation models*, which was introduced in [5] as a framework for binary merging in which the merging is achieved via a negotiation-like process. We extend this framework so that it can handle information coming from n sources for $n \in \mathbb{N}$, and show how a given belief negotiation model yields an SC function.

The plan of the paper is as follows. We begin in Section 2 by formally defining SC functions via a small list of basic properties we expect such an operation to satisfy. We show how one of these basic properties allows us to derive, from a given *social* contraction function, a list of *individual* contraction functions (in the traditional belief revision sense as described above)—one for each information source. We also describe how a given SC function yields a merging operator via a kind of "generalised" Levi Identity before ending the section with a look at a few possible additional postulates for social contraction, relating to the idea—familiar from belief revision—of *minimal change*. The rest of the paper is devoted to belief negotiation. The extended framework is set down in Section 3, where it is shown how each (extended) belief negotiation model yields an SC function and, conversely, how every SC function can be said to arise in this way. As we will see, the framework is set at a very abstract level. Section 4 is all about putting a little more flesh on the bones. Making heavy use of Spohn-type rankings [22] we provide two, intuitively plausible, instantiations of the parameters of a belief negotiation model, giving in the process two concrete families of SC functions. We characterise the behaviour of the individual contraction functions as well as the merging operators which are derivable from these particular families. It turns out that they are all familiar from the literature. We thus give a new angle on these operators by providing new "negotiation-style" characterisations for them. We also test the SC functions from each of these two families against the extra minimal change postulates from Section 2. We will see that the SC functions from the second family fare better than those from the first in this regard. We conclude in Section 5. Proofs of our results are contained in Appendix A.

1.1. Preliminaries

In this paper we shall follow the example of the papers on merging mentioned at the start of the introduction, and assume a very simple propositional setting for the merging problem. (For more complex settings, e.g., where the items to be merged consist of formulas of first-order logic, or settings from the area of database theory, we refer the reader to, e.g., [2–4,9,10].) We let $\mathscr{W}$ be the (finite) set of worlds, i.e., truth-assignments, associated with some fixed background propositional language generated from finitely many propositional variables. The set of all non-empty subsets of $\mathscr{W}$ we denote by $\mathscr{B}$. Given $S \subseteq \mathscr{W}$, we use $\overline{S}$ to denote $\mathscr{W} - S$. We assume throughout that we have a fixed finite set $Sources = \{0, 1, \ldots, n\}$ of information sources ($n \geqslant 1$). We work semantically throughout, so each item of information provided by a source i will take the form of a set $S_i \in \mathscr{B}$ (so no source ever provides the "inconsistent" information $\emptyset$). Such an S_i should be interpreted as the information that the actual "true" world is one of the worlds in S_i. An *information profile* (relative to *Sources*) is an element of $\mathscr{B}^{Sources}$, i.e., a particular assignment of elements of $\mathscr{B}$ to the sources. We shall use vector notation $\vec{S}$, $\vec{S}^1$, etc. to denote information profiles, with $\vec{S} = (S_0, S_1, \ldots, S_n)$, $\vec{S}^1 = (S^1_0, S^1_1, \ldots, S^1_n)$, etc. The idea is that S_i is the information in $\vec{S}$ belonging to source i. We will say that an information profile $\vec{S}$ is *consistent* when $\cap_i S_i \neq \emptyset$, otherwise it is *inconsistent*. Given two information profiles $\vec{S}^1$ and $\vec{S}^2$, we will write $\vec{S}^1 \subseteq \vec{S}^2$ to mean $S^1_i \subseteq S^2_i$ for all $i \in Sources$. If $\vec{S}^1 \subseteq \vec{S}^2$ and $\vec{S}^2 \nsubseteq \vec{S}^1$ then we will write $\vec{S}^1 \subset \vec{S}^2$. Finally if $\mathbf{f}$ is a function with codomain $\mathscr{B}^{Sources}$, we will use $\mathbf{f}_i(\vec{S})$ to denote the $i+1$th element of $\mathbf{f}(\vec{S})$, i.e., we will have $\mathbf{f}(\vec{S}) = (\mathbf{f}_0(\vec{S}), \mathbf{f}_1(\vec{S}), \ldots, \mathbf{f}_n(\vec{S}))$.

2. Social contraction functions

Our first aim is to get a formal definition of SC functions up and running. Intuitively we want an SC function to be a function $\mathbf{f} : \mathscr{B}^{Sources} \rightarrow \mathscr{B}^{Sources}$ which, given an information profile $\vec{S}$ provided by *Sources*, returns a new information profile $\mathbf{f}(\vec{S})$ which represents $\vec{S}$ modified so that its entries are jointly consistent. We immediately require the following three basic properties of such an $\mathbf{f}$:

(sc1) $\vec{S} \subseteq \mathbf{f}(\vec{S})$.
(sc2) $\mathbf{f}(\vec{S})$ is consistent.
(sc3) If $\vec{S}$ is consistent then $\mathbf{f}(\vec{S}) = \vec{S}$.

Rule (sc1) decrees that the modification is carried out by *weakening* the individual items of information. Hence, to obtain consistency, we require that some information may be *taken away* from the original items S_i. However, no information is allowed to be *added*. (This justifies the name "social *contraction*".)[1] Rule (sc2) says that the end results of all these weakenings should be jointly consistent. Rule (sc3) says that if $\vec{S}$ is already consistent then

[1] But see [6] for a treatment of (individual, not social) information-removal operators in which the removal of a piece of information *can* directly lead to the introduction of *new* information.

no modification is necessary. In addition to these three properties, we shall also find it convenient to assume that, amongst the sources, there is one distinguished source who is *completely reliable*, in the sense that any information provided by this source can safely be assumed to be true and so should never be weakened. *We fix source 0 to be this completely reliable source*, and reflect this by insisting on the following rule for SC functions:

(sc4) $\mathbf{f}_0(\vec{S}) = S_0$.

We will denote the set of sources *minus* 0 by $Sources^+$. We recognise that the existence of such a completely reliable source is not guaranteed in practice. However, situations where it is absent can be modelled by simply taking S_0 to be the "trivial" information, i.e., take $S_0 = \mathscr{W}$. We now make the following definition.

Definition 1. Let $\mathbf{f} : \mathscr{B}^{Sources} \to \mathscr{B}^{Sources}$ be a function. Then $\mathbf{f}$ is a *social contraction function* (*relative to Sources*) iff it satisfies (sc1)–(sc4).

We now give a couple of simple examples of SC functions.

Example 2. (i) A very simple example of an SC function is the "trivial" SC function $\mathbf{f}^{\text{triv}}$ which, given an input information profile $\vec{S}$, just returns $\vec{S}$ if this is consistent, and which otherwise weakens all items of information (except the completely reliable S_0) right out to $\mathscr{W}$. Precisely, for each $i \in Sources$,

$$\mathbf{f}_i^{\text{triv}}(\vec{S}) = \begin{cases} S_i & \text{if } i = 0 \text{ or } \vec{S} \text{ is consistent} \\ \mathscr{W} & \text{otherwise.} \end{cases}$$

According to this operator, *all* items of information (except S_0) are effectively discarded as soon as input $\vec{S}$ is inconsistent. This marks $\mathbf{f}^{\text{triv}}$ down as quite a "wasteful" operator.

(ii) A slightly more refined version of this is the SC function for which $\mathbf{f}(\vec{S})$ again returns $\vec{S}$ if this is consistent, and otherwise weakens each S_i by just adding S_0:

$$\mathbf{f}_i(\vec{S}) = \begin{cases} S_i & \text{if } \vec{S} \text{ is consistent} \\ S_i \cup S_0 & \text{otherwise.} \end{cases}$$

It is easy to verify that both the above functions satisfy (sc1)–(sc4). Some more sophisticated examples of SC functions will be presented in Section 4, after the ideas of belief negotiation have been introduced.

A benefit of including (sc4) among our basic postulates is that it gives us access to a list of individual, "local" contraction functions — one for each $i \in Sources^+$. These functions reveal, for each source i, how any item of information from i would be weakened in the face of a single second item which is considered completely reliable.

Definition 3. Let $\mathbf{f}$ be an SC function and let $i \in Sources^+$. We define the function $\ominus_i^{\mathbf{f}} : \mathscr{B} \times \mathscr{B} \to \mathscr{B}$ by, for all $S, T \in \mathscr{B}$, $S \ominus_i^{\mathbf{f}} T = \mathbf{f}_i(\vec{U})$, where $\vec{U} \in \mathscr{B}^{Sources}$ is such that $U_i = S$, $U_0 = T$ and $U_j = \mathscr{W}$ for all $j \notin \{0, i\}$. We call $\ominus_i^{\mathbf{f}}$ *i's individual contraction function* (*relative to* $\mathbf{f}$).

(E.g., if $n = 3$, then $S \ominus_2^{\mathbf{f}} T$ is the 3rd entry of the 4-tuple $\mathbf{f}(T, \mathscr{W}, S, \mathscr{W})$.) Thus $S \ominus_i^{\mathbf{f}} T$ represents the result—according to $\mathbf{f}$—of weakening information S from source i so that it becomes consistent with T. We have the following proposition.

Proposition 4. *Let* $\mathbf{f}$ *be an SC function and let* $i \in Sources^+$. *Then* $\ominus_i^{\mathbf{f}}$ *satisfies*
(ind1) $S \subseteq S \ominus_i^{\mathbf{f}} T$
(ind2) $(S \ominus_i^{\mathbf{f}} T) \cap T \neq \emptyset$
(ind3) *If* $S \cap T \neq \emptyset$ *then* $S \ominus_i^{\mathbf{f}} T = S$

The properties (ind1)–(ind3) essentially correspond to the well-known basic AGM postulates for contraction (1) *minus* the Recovery postulate, which in our notation would correspond to "$S \ominus_i^{\mathbf{f}} T \subseteq S \cup T$". It will become apparent in Section 4 that the $\ominus_i^{\mathbf{f}}$ do not generally satisfy this much debated (see [11, pp. 71–74]) property.

Recall that a principle motivating factor behind defining SC functions was to use them as a stepping-stone to defining merging operators. Under this view, the result of the SC operation on $\vec{S}$ represents an *intermediate* stage in the merging of the information items in $\vec{S}$, in which simple conjunction of the information items can then be easily facilitated. From a given SC function $\mathbf{f}$, we define the merging operator $\Delta_{\mathbf{f}}$ relative to *Sources* using a kind of "generalised" Levi Identity. We set, for each information profile $\vec{S}$,

$$\Delta_{\mathbf{f}}(\vec{S}) = \bigcap_{i=0}^{n} \mathbf{f}_i(\vec{S}).$$

Our basic postulates for $\mathbf{f}$ immediately yield a corresponding set of basic properties for $\Delta_{\mathbf{f}}$: (sc2) gives $\Delta_{\mathbf{f}}(\vec{S}) \neq \emptyset$, while from (sc3) we get that $\vec{S}$ is consistent implies $\Delta_{\mathbf{f}}(\vec{S}) = \bigcap_i S_i$. Meanwhile (sc4) gives us $\Delta_{\mathbf{f}}(\vec{S}) \subseteq S_0$, i.e., the result of the merging must always imply the information provided by source 0. In this respect $\Delta_{\mathbf{f}}$ resembles what is referred to by Konieczny and Pino-Pérez as a *merging operator with integrity constraints*, or IC merging operator for short [13], S_0 here taking the role of the integrity constraints in their framework.[2] (For a more complicated treatment of integrity constraints, see [3].)

[2] At this point it is natural to ask whether it is possible to take the converse direction and derive an SC function *from* a given IC merging operator, just like, in belief revision, it is possible to derive a contraction operator from a given revision operator via the Harper Identity [8]. This question will be taken up in future work.

2.1. More postulates: minimal change

The postulates (sc1)–(sc4) form our core set of postulates for SC functions, but there is clearly scope for other desirable properties to be put forward. One possible source for such further postulates is the idea of *minimal change*, i.e., the idea that the modification of $\vec{S}$ to achieve consistency should be kept as "small" as possible.[3] Our condition (sc3) can already be said to be a mild embodiment of this idea. In this subsection we look at a couple of ways in which it can be taken further. The first rule we consider is the following:

(sc5) For all $i \in Sources^+$, if $S_i \cap \bigcap_{j \neq i} \mathbf{f}_j(\vec{S}) \neq \emptyset$ then $\mathbf{f}_i(\vec{S}) = S_i$.

The motivation behind this rule is the feeling that, for each $i \in Sources^+$, we should take $\mathbf{f}_i(\vec{S}) = S_i$ whenever possible. (Recall we already have $\mathbf{f}_0(\vec{S}) = S_0$ by (sc4).) Clearly if $S_i \cap \bigcap_{j \neq i} \mathbf{f}_j(\vec{S}) \neq \emptyset$ then it is possible. It is easy to see that, in the presence of (sc1) and (sc4), (sc5) implies (sc3):

Proposition 5. *Let* $\mathbf{f} : \mathscr{B}^{Sources} \to \mathscr{B}^{Sources}$ *be a function which satisfies* (*sc*1), (*sc*4) *and* (*sc*5). *Then* $\mathbf{f}$ *satisfies* (*sc*3).

It is also quite easy to see that the trivial SC function $\mathbf{f}^{\text{triv}}$ from Example 2 does not satisfy (sc5). Hence (sc5) doesn't hold in general for SC functions. However, even though (sc5) may be appealing from a minimal change point of view, its adoption can lead to counter-intuitive results, as the following example shows.

Example 6. Suppose we have three sources, i.e., $n = 2$. Suppose source 1 provides the information $S \neq \mathscr{W}$, source 2 provides the complete opposite information $\overline{S}$, and the completely reliable source 0 provides only the trivial information $\mathscr{W}$. We first claim that for any SC function $\mathbf{f}$ relative to these sources which satisfies (sc5) we have either $\mathbf{f}_1(\mathscr{W}, S, \overline{S}) = S$ or $\mathbf{f}_2(\mathscr{W}, S, \overline{S}) = \overline{S}$. To see this, suppose $\mathbf{f}_1(\mathscr{W}, S, \overline{S}) \neq S$. Then, by (sc5), we must have $S \cap \mathbf{f}_0(\mathscr{W}, S, \overline{S}) \cap \mathbf{f}_2(\mathscr{W}, S, \overline{S}) = \emptyset$. Now we know by (sc4) (or (sc1)) that $\mathbf{f}_0(\mathscr{W}, S, \overline{S}) = \mathscr{W}$. Hence we have $S \cap \mathbf{f}_2(\mathscr{W}, S, \overline{S}) = \emptyset$, i.e., $\mathbf{f}_2(\mathscr{W}, S, \overline{S}) \subseteq \overline{S}$. Since we also have $\overline{S} \subseteq \mathbf{f}_2(\mathscr{W}, S, \overline{S})$ by (sc1), we conclude that $\mathbf{f}_2(\mathscr{W}, S, \overline{S}) = \overline{S}$ which proves the claim. Given this, we have for the corresponding merging operator that either $\Delta_{\mathbf{f}}(\mathscr{W}, S, \overline{S}) \subseteq S$ or $\Delta_{\mathbf{f}}(\mathscr{W}, S, \overline{S}) \subseteq \overline{S}$. Hence when merging S and $\overline{S}$ we are forced to accept one or the other. However one can easily imagine a situation where we are unable to find any reason to prefer S to $\overline{S}$ or vice-versa (e.g. sources 1 and 2 are equally reliable, equally convinced their information is correct etc.). In this case it would not seem irrational to withhold judgement on whether S or $\overline{S}$ holds in the merging and to expect, say, $\Delta_{\mathbf{f}}(\mathscr{W}, S, \overline{S}) = \mathscr{W}$. Merging using an SC function which satisfies (sc5) rules out this possibility.

This is reminiscent of the problems with so-called *maxichoice* contraction and revision in the belief change literature (see [11, pp. 76–77, 209–210]). To understand why, it is helpful to change perspective slightly. For each SC function $\mathbf{f}$ and each information profile $\vec{S}$ define the set $X_{\mathbf{f}}(\vec{S}) \subseteq Sources^+$ by

$$X_{\mathbf{f}}(\vec{S}) = \{i \in Sources^+ \mid \mathbf{f}_i(\vec{S}) = S_i\}.$$

In other words, given that *Sources* provides the information $\vec{S}$, $X_{\mathbf{f}}(\vec{S})$ is the set of sources (other than 0) who do not weaken their information according to $\mathbf{f}$. The principle of minimal change suggests we should take $X_{\mathbf{f}}(\vec{S})$ to be an inclusion-maximal subset of $Sources^+$. This is ensured by the following rule, which bears a strong resemblance to the contraction postulate "Fullness" [11, p. 77] which, in turn, is a characteristic postulate of maxichoice contraction:

(sc5+) For all $i \in Sources^+$, if $S_i \cap \left(\bigcap_{j \in X_{\mathbf{f}}(\vec{S})} S_j\right) \cap S_0 \neq \emptyset$ then $i \in X_{\mathbf{f}}(\vec{S})$.

As the next proposition shows, in the presence of (sc4), (sc5+) implies (sc5). However, in the additional presence of the following strengthening of (sc1), (sc5) becomes *equivalent* to (sc5+):

(sc1+) For all $i \in Sources$, either $\mathbf{f}_i(\vec{S}) = S_i$ or $\mathbf{f}_i(\vec{S}) = \mathscr{W}$.

Proposition 7. *Let* $\mathbf{f} : \mathscr{B}^{Sources} \to \mathscr{B}^{Sources}$ *be a function which satisfies* (*sc*4). *Then, if* $\mathbf{f}$ *satisfies* (*sc*5+), *then* $\mathbf{f}$ *satisfies* (*sc*5). *Furthermore, if* $\mathbf{f}$ *additionally satisfies* (*sc*1+) *then the converse holds.*

The rule (sc1+) says, in effect, that the information from each source is either kept or discarded completely.[4]

Although Example 6 suggests (sc5) may be too strong for SC functions, possible weakenings of it are at hand. One, which brings the individual contraction functions into the picture, is the following:

(sc6) For all $i \in Sources^+$, if $S_i \cap \bigcap_{j \neq i} \mathbf{f}_j(\vec{S}) \neq \emptyset$ then $\mathbf{f}_i(\vec{S}) \subseteq S_i \ominus_i^{\mathbf{f}} \overline{S_i}$.

[3] The idea of minimal change is also a major consideration in several other merging formalisms such as those presented in [2,3,10].

[4] Precisely such an assumption is made explicitly in [7]. Its adoption here would effectively reduce social contraction to something akin to *belief base contraction* [11].

Note that $S_i \ominus_i^{\mathbf{f}} \overline{S_i}$ is the result of weakening S_i so that it becomes consistent with $\overline{S_i}$ and so, intuitively, contains those worlds in $\overline{S_i}$ which, at least from i's viewpoint, are considered the most plausible. Hence the consequent of (sc6) essentially says that if $\mathbf{f}_i(\vec{S})$ *has* to contain worlds outside of S_i, then it should contain only the most plausible ones. Unfortunately this rule is rather too weak to enforce minimal change, a fact which can be seen by noting that the trivial SC function $\mathbf{f}^{\mathrm{triv}}$ validates it. Rather, (sc6) represents some sort of "coherence" condition on the results of performing social contraction on different, but related information profiles (in this case $\vec{S}$ and $(\overline{S_i}, \mathscr{W}, \ldots, S_i, \ldots, \mathscr{W})$). Another weakening of (sc5) is the following:

(sc7) For all $i \in Sources^{+}$, if $\mathbf{f}_i(\vec{S}) \neq S_i$ then there exists some consistent $\vec{T}$ such that $\vec{S} \subseteq \vec{T} \subseteq \mathbf{f}(\vec{S})$ and $S_i \cap \bigcap_{j \neq i} T_j = \emptyset$.

This rule (which has a similar form to the postulate "Relevance" from belief base contraction [11, p. 68]) can be explained as follows: If, for every consistent information profile $\vec{T}$ lying "between" $\vec{S}$ and $\mathbf{f}(\vec{S})$, it is possible to reduce T_i to S_i without incurring inconsistency, then it seems safe to say that S_i does not in any way contribute to any inconsistency arising in $\vec{S}$. Hence (sc7) provides a way of saying that source i's information is weakened only if it somehow contributes to the inconsistency of the information profile $\vec{S}$. Although weaker than (sc5), (sc7) still manages to be stronger than (sc3) (with the help once again of (sc4)):

Proposition 8. *Let* $\mathbf{f} : \mathscr{B}^{Sources} \to \mathscr{B}^{Sources}$ *be a function which satisfies* (*sc*4) *and* (*sc*7). *Then* $\mathbf{f}$ *satisfies* (*sc*3).

Meanwhile, unlike (sc6), (sc7) still manages to be strong enough to exclude $\mathbf{f}^{\mathrm{triv}}$, as the following example shows.

Example 9. Assume $Sources = \{0,1,2\}$ and that $\vec{S} = (S, \overline{S}, S)$, where $S \in \mathscr{B}$ is such that $S \neq \mathscr{W}$. Then since $\vec{S}$ is inconsistent we have $\mathbf{f}^{\mathrm{triv}}(\vec{S}) = (S, \mathscr{W}, \mathscr{W})$. Hence we see $S_2 = S \neq \mathscr{W} = \mathbf{f}_2^{\mathrm{triv}}(\vec{S})$. If $\mathbf{f}^{\mathrm{triv}}$ satisfied (sc7) we would deduce the existence of some consistent $\vec{T} \in \mathscr{B}^{\{0,1,2\}}$ such that $\vec{S} \subseteq \vec{T} \subseteq \mathbf{f}^{\mathrm{triv}}(\vec{S})$ and $S_2 \cap T_0 \cap T_1 = \emptyset$. Since both $S_2 = S$ and $T_0 = S$ (this latter holding since $S_0 \subseteq T_0 \subseteq \mathbf{f}_0^{\mathrm{triv}}(\vec{S})$ and $\mathbf{f}_0^{\mathrm{triv}}(\vec{S}) = S_0$, hence $T_0 = S_0 = S$), we deduce from $S_2 \cap T_0 \cap T_1 = \emptyset$ that $T_1 \subseteq \overline{S}$. But using this with the fact that $T_0 = S$ gives us $\vec{T}$ is inconsistent—contradiction. Hence it cannot be that $\mathbf{f}^{\mathrm{triv}}$ satisfies (sc7).

Our final postulate is motivated by the feeling that social contraction should be entirely expressible in terms of the individual contraction functions.

(sc8) For all $i \in Sources^{+}$, $\mathbf{f}_i(\vec{S}) = S_i \ominus_i^{\mathbf{f}} \left(\bigcap_{j \neq i} \mathbf{f}_j(\vec{S}) \right)$.

This postulate can also be interpreted as saying that the outcome $\mathbf{f}(\vec{S})$ of an operation of social contraction represents a kind of *equilibrium* state. One in which each source's information S_i is weakened just enough—according to that source's own individual contraction function—to be consistent with the joint result of the weakenings of all the other sources. Since, by Proposition 4, $\ominus_i^{\mathbf{f}}$ satisfies (ind3), it is easy to see that any SC function satisfying (sc8) also satisfies (sc5). In fact, as the following result confirms, only the "$\subseteq$" direction of (sc8) is needed to prove (sc5).

Proposition 10. *Let* $\mathbf{f} : \mathscr{B}^{Sources} \to \mathscr{B}^{Sources}$ *be an SC function such that, for all* $i \in Sources^{+}$ *and all* $\vec{S} \in \mathscr{B}^{Sources}$, $\mathbf{f}_i(\vec{S}) \subseteq S_i \ominus_i^{\mathbf{f}} \left(\bigcap_{j \neq i} \mathbf{f}_j(\vec{S}) \right)$. *Then* $\mathbf{f}$ *satisfies* (*sc*5).

3. Extended belief negotiation models

So far we have examined social contraction from a strictly postulational viewpoint. In the rest of the paper we adopt another, more procedural, perspective. In [5] the framework of *belief negotiation models* was introduced as a framework for merging together information from just *two* different sources. The idea was that the pieces of information were weakened *incrementally* via a negotiation-like process until "common ground" was reached, i.e., until they became consistent with one another. The purpose of this section is to extend this framework so that it handles information coming from $n + 1$ different sources (one of which is considered completely reliable) and show how each such extended belief negotiation model $\mathscr{N}$ yields an SC function $\mathbf{f}^{\mathscr{N}}$. Let us begin with a rough description of the framework.[5]

Suppose the information profile $\vec{S}$ is provided by *Sources*. The idea is that we determine $\mathbf{f}^{\mathscr{N}}(\vec{S})$ as follows. We start off with the information profile $\vec{S}^0 = \vec{S}$. If $\vec{S}^0$ is consistent then we just take $\mathbf{f}^{\mathscr{N}}(\vec{S}) = \vec{S}^0$. But if $\vec{S}^0$ is inconsistent then we perform what may be thought of as a "round of negotiation" which is just a contest between the sources. The losers of this contest (for there may be several) must then "make some concessions", i.e., make some weakening of their position by admitting more possibilities, while the others stay the same. Thus we arrive at the new information profile $\vec{S}^1$ where $\vec{S}^0 \subseteq \vec{S}^1$. Now if $\vec{S}^1$ is consistent then we set $\mathbf{f}^{\mathscr{N}}(\vec{S}) = \vec{S}^1$.

[5] We remark that this framework shares some similarities with the abstract formalisation of negotiation found in [25]. For another recent attempt at bringing ideas from belief revision and negotiation together see [18]. A more detailed treatment of the subject of negotiation can be found in [24].

Otherwise the next round of negotiation takes place. Once again the losers of this round make concessions, and we keep going like this until $\vec{S}^j$ is consistent, at which point we set $\mathbf{f}^{\mathcal{N}}(\vec{S}) = \vec{S}^j$. Now let us spell this out in detail.

Let Ω denote the set of all finite sequences of information profiles. Given $\omega = (\vec{S}^0, \ldots, \vec{S}^m) \in \Omega$ we will say that ω is *increasing* iff $\vec{S}^j \subseteq \vec{S}^{j+1}$ for all $j = 0, 1, \ldots, m-1$. We define the set of sequences $\Sigma \subseteq \Omega$ by

$$\Sigma = \{\omega = (\vec{S}^0, \ldots, \vec{S}^m) \in \Omega \mid \omega \text{ is increasing, and } \vec{S}^m \text{ is inconsistent}\}.$$

A sequence $\sigma = (\vec{S}^0, \ldots, \vec{S}^m) \in \Sigma$ represents a possible stage in the unfinished (since $\vec{S}^m$ is inconsistent) negotiation process starting with $\vec{S}^0$. Here, the information profile $\vec{S}^m$ describes the *current standpoints* of the sources at stage σ. Given $j < m$, we let σ_j denote that sequence consisting of the first $j+1$ entries in σ, i.e., $\sigma_j = (\vec{S}^0, \ldots, \vec{S}^j)$.

In the simple negotiation scenario described above there were two ingredients in need of further specification. Firstly, we need to know how a round of negotiation is carried out. To begin with, we don't worry about the precise details. We simply assume the existence of a function $g : \Sigma \to 2^{Sources^+}$ which selects, at each negotiation stage σ, which parties should make concessions. In other words g returns the losers of the negotiation round at stage σ. Note that here we are building in our assumption that source 0 is completely reliable (and so never loses a round) by taking the codomain of g to be $2^{Sources^+}$ rather than $2^{Sources}$. We make two more mild restrictions on g. First, in order to avoid deadlock we need to assume that at least one party must weaken at each stage:

(g0a) $g(\sigma) \neq \emptyset$.

Second, suppose we reach a negotiation stage $\sigma = (\vec{S}^0, \ldots, \vec{S}^m)$ such that $S_i^m = \mathcal{W}$ for some $i \in Sources^+$. Then obviously at this stage i's information cannot be weakened any further. We restrict g so that it selects only sources who still have "room to manoeuvre".

(g0b) $i \in g(\sigma)$ implies $S_i^m \neq \mathcal{W}$ (where $\sigma = (\vec{S}^0, \ldots, \vec{S}^m)$).

The second missing ingredient is then to decide what concessions the losers of a negotiation round should make. Once again we initially abstract away from the actual process used to determine this and assume only that we are given, for each $\sigma = (\vec{S}^0, \ldots, \vec{S}^m) \in \Sigma$, a function $\blacktriangledown_\sigma : Sources^+ \to \mathcal{B}$ with the interpretation that $\blacktriangledown_\sigma(i)$ represents the weakening of S_i^m that would be made, *given that i were chosen to weaken at stage* σ. Once again to avoid deadlock, we require that this weakening be strict, unless of course $S_i^m = \mathcal{W}$:

($\blacktriangledown$0a) $S_i^m \subseteq \blacktriangledown_\sigma(i)$
($\blacktriangledown$0b) $\blacktriangledown_\sigma(i) = S_i^m$ implies $S_i^m = \mathcal{W}$.

The reader may notice that, even though we are requiring that $\blacktriangledown_\sigma(i)$ be a strict weakening of S_i^m for *all* $i \in Sources^+$, these weakenings will only actually be "carried out" if i is a loser of the negotiation round at stage σ, i.e., $i \in g(\sigma)$. Hence to avoid deadlock it is really only necessary that $\blacktriangledown_\sigma(i)$ be a strict weakening of S_i^m **for some** $i \in g(\sigma)$. Our stronger requirement above comes from our desire to keep our conditions on the $\blacktriangledown_\sigma(i)$ independent from our conditions on g. Note also that here we again identify information removal with information weakening. We could, for a more general treatment, weaken these properties on the $\blacktriangledown_\sigma$, although then, of course, termination of the negotiation process would no longer be guaranteed.

We can now make the following definition.

Definition 11. An *extended belief negotiation model* (*relative to Sources*) is a pair $\mathcal{N} = \langle g, \{\blacktriangledown_\sigma\}_{\sigma\in\Sigma}\rangle$ where $g : \Sigma \to 2^{Sources^+}$ is a function which satisfies (g0a) and (g0b), and, for each $\sigma \in \Sigma$, $\blacktriangledown_\sigma : Sources^+ \to \mathcal{B}$ is a function which satisfies ($\blacktriangledown$0a) and ($\blacktriangledown$0b).

From now on when we write "belief negotiation model" we will mean an extended belief negotiation model in the sense of the above definition.[6]

Example 12. (i) Perhaps the simplest example of a belief negotiation model is $\mathcal{N}^{\text{triv}} = \langle g, \{\blacktriangledown_\sigma\}_{\sigma\in\Sigma}\rangle$ where we take $g(\sigma) = Sources^+$ and $\blacktriangledown_\sigma(i) = \mathcal{W}$ for all $\sigma \in \Sigma$ and $i \in Sources^+$.

(ii) Another possibility for g would be to select at stage $\sigma = (\vec{S}^0, \ldots, \vec{S}^m) \in \Sigma$ all sources whose current standpoint is not implied by the information of source 0, i.e., take $g(\sigma) = \{i \in Sources^+ \mid S_0^m \not\subseteq S_i^m\}$. Another possibility for the $\blacktriangledown_\sigma$ would be to add all of S_0^m to S_i^m if this produces a strict weakening, otherwise to just add all worlds, i.e., for all $\sigma = (\vec{S}^0, \ldots, \vec{S}^m) \in \Sigma$ and $i \in Sources^+$,

$$\blacktriangledown_\sigma(i) = \begin{cases} S_i^m \cup S_0^m & \text{if } S_0^m \not\subseteq S_i^m \\ \mathcal{W} & \text{otherwise.} \end{cases}$$

We will give some more sophisticated examples of belief negotiation models in Section 4.

---

[6] There are a couple of slight notational differences between this paper and [5]. In the latter paper the function g picked up the actual information items to be weakened rather than naming the sources from which they came. Similarly the functions $\blacktriangledown_\sigma$ were defined directly on the elements of S_i^m rather than the set of sources.

A belief negotiation model $\mathcal{N}$ then uniquely determines, for any given information profile $\vec{S}$ provided by *Sources*, the complete process of negotiation on $\vec{S}$. This process is returned by the function $f^{\mathcal{N}} : \mathscr{B}^{Sources} \to \Omega$ given by

$$f^{\mathcal{N}}(\vec{S}) = \sigma = (\vec{S}^0, \ldots, \vec{S}^k)$$

where (i) $\vec{S}^0 = \vec{S}$, (ii) k is minimal such that $\vec{S}^k$ is consistent, and (iii) for each $0 \leqslant j < k$ we have, for each $i \in Sources$,

$$S_i^{j+1} = \begin{cases} \blacktriangledown_{\sigma_j}(i) & \text{if } i \in g(\sigma_j) \\ S_i^j & \text{otherwise.} \end{cases}$$

It should be clear that the restrictions we have placed on g and the $\blacktriangledown_\sigma$ (along with our assumption that $\mathscr{W}$ is finite) guarantee the existence of the minimal k in (ii) above. A belief negotiation model $\mathcal{N}$ thus yields a function $\mathbf{f}^{\mathcal{N}} : \mathscr{B}^{Sources} \to \mathscr{B}^{Sources}$, via $f^{\mathcal{N}}$ above, by simply taking $\mathbf{f}^{\mathcal{N}}(\vec{S}) = \vec{S}^k$. It is straightforward to check that $\mathbf{f}^{\mathcal{N}}$ forms an SC function. Furthermore, in fact *every* SC function can be said to arise in this way.

Theorem 13. *Let* $\mathbf{f} : \mathscr{B}^{Sources} \to \mathscr{B}^{Sources}$ *be a function. Then* $\mathbf{f}$ *is an SC function iff* $\mathbf{f} = \mathbf{f}^{\mathcal{N}}$ *for some belief negotiation model* $\mathcal{N}$.

The reader may like to verify that the function $\mathbf{f}^{\mathcal{N}^{\text{triv}}}$ generated from the belief negotiation model $\mathcal{N}^{\text{triv}}$ from Example 12(i) is in fact equal to the trivial SC function $\mathbf{f}^{\text{triv}}$ from Example 2(i), while the function $\mathbf{f}^{\mathcal{N}}$ generated from the belief negotiation model $\mathcal{N}$ from Example 12(ii) is equal to the SC function given in Example 2(ii).

In what follows we use $\Delta_{\mathcal{N}}$ to denote the merging operator defined from $\mathbf{f}^{\mathcal{N}}$, and $\ominus_i^{\mathcal{N}}$ to denote source i's individual contraction function $\ominus_i^{\mathbf{f}^{\mathcal{N}}}$ relative to $\mathbf{f}^{\mathcal{N}}$. A point to note about these latter functions is that they depend only on the functions $\blacktriangledown_\sigma$, i.e., we have the following result.

Proposition 14. *Let* $\mathcal{N} = \langle g, \{\blacktriangledown_\sigma\}_{\sigma\in\Sigma}\rangle$ *and* $\mathcal{N}' = \langle g', \{\blacktriangledown_\sigma\}_{\sigma\in\Sigma}\rangle$ *be two belief negotiation models which differ only on their first component. Then, for each* $i \in Sources^+$, *we have* $\ominus_i^{\mathcal{N}} = \ominus_i^{\mathcal{N}'}$.

4. Instantiating the framework

A natural question to ask about the preceding framework is: where do the functions g and $\blacktriangledown_\sigma$ of a belief negotiation model come from? In this section we explore some possibilities—one for the $\blacktriangledown_\sigma$ and two for g, leading to two different concrete families of SC functions. To help us do this we first need to make some extra demands on the type of information provided by our sources. We assume that each source $i \in Sources^+$ provides not only a set $S_i \in \mathscr{B}$, but also some indication of the *plausibility* of *all* the worlds in $\mathscr{W}$. Such an indication is provided by a *ranking*.

Definition 15. A *ranking* is a function $r : \mathscr{W} \to \mathbb{N}$. We extend such an r to a function on $\mathscr{B}$ by setting, for each $T \in \mathscr{B}$, $r(T) = \min_{w\in T} r(w)$. Given $S \in \mathscr{B}$ we say that r is a ranking *anchored on* S iff $r^{-1}(0) = S$.

Example 16. To give an example of a ranking, let's assume our background propositional language contains just two propositional variables, leading $\mathscr{W}$ to contain just four worlds which we denote here by a, b, c, d. Then we can specify the ranking r in tabular form as follows:

	0	1	2	3
r	a,b		c	d

Here, the columns correspond to ranks, so in fact we have $r(a) = r(b) = 0$, $r(c) = 2$ and $r(d) = 3$. We also have $r(\{c,d\}) = \min\{r(c), r(d)\} = 2$ and $r(\{a,c,d\}) = \min\{r(a), r(c), r(d)\} = 0$. Meanwhile, since $r^{-1}(0) = \{a,b\}$, r is anchored on $\{a,b\}$.

Such rankings, or variants thereof, are a popular tool in knowledge representation. They can be traced back to the work of [22] and indeed have already been employed in the context of both merging (see e.g. [17,19]) and belief revision (see e.g. [23]). A ranking provides, for each $w \in \mathscr{W}$, a measure of the plausibility of w being the actual world. The lower $r(w)$ is, the more plausible it is considered to be. The plausibility $r(T)$ of a *set* T of worlds is identified with that of the most plausible worlds in T. Rankings also allow us to talk about *degrees of certainty* or *degrees of belief*. Given $S \in \mathscr{B}$, we can interpret $r(\overline{S})$ as the degree of certainty that the world is in S—the higher $r(\overline{S})$ is, i.e., the more *implausible* $\overline{S}$ is, the more certain it is that S contains the actual world. We now assume that each time a source $i \in Sources^+$ provides the information S_i, he provides along with it a ranking anchored on S_i. Formally, we assume we are given a *ranking assignment* for *Sources*.

Definition 17. A *ranking assignment* (relative to *Sources*) is a function R which assigns, to each $i \in Sources^+$ and $S \in \mathscr{B}$, a ranking $[R_i(S)]$ anchored on S.

Note we assume source 0 does not provide a ranking, just S_0 as normal. We also make an assumption of *commensurability* [19], i.e., that all sources use the same *scale* when ranking the worlds according to plausibility. Given this definition, we are now in a position to describe our first instantiation of the framework.

4.1. First instantiation

How can we use a ranking assignment R to define suitable functions g and $\blacktriangledown_\sigma$? Turning first to g, our idea is this: the losers of the negotiation round at stage $\sigma = (\vec{S}^0, \ldots, \vec{S}^m)$ should be those sources i who are the *least certain* about their current standpoint S_i^m, according to the ranking $[R_i(S_i^0)]$ which they have provided along with their initial information S_i^0. Recall that the *lower* the number $[R_i(S_i^0)](\overline{S_i^m})$ is, the *less certain* i is about S_i^m. Thus, precisely, we define g_1 from R by setting

$$g_1(\sigma) = \{i \in Sources^+ \mid S_i^m \neq \mathscr{W} \text{ and } [R_i(S_i^0)](\overline{S_i^m}) \leqslant [R_j(S_j^0)](\overline{S_j^m}) \text{ for all } j \text{ such that } S_j^m \neq \mathscr{W}\}.$$

As for defining $\blacktriangledown_\sigma$, the method we choose is quite simple. We assume that, for each $\sigma = (\vec{S}^0, \ldots, \vec{S}^m) \in \Sigma$, if source i has to weaken at stage σ, he does so by adding to S_i^m those worlds not already in S_i^m which are the most plausible according to the ranking i has provided with his initial information S_i^0. More precisely we set

$$\blacktriangledown_\sigma(i) = S_i^m \cup \{w \in \overline{S_i^m} \mid [R_i(S_i^0)](w) \leqslant [R_i(S_i^0)](w') \text{ for all } w' \in \overline{S_i^m}\}.$$

Given a ranking assignment R, we let $\mathscr{N}_1(R)$ denote the belief negotiation model $\langle g_1, \{\blacktriangledown_\sigma\}_{\sigma\in\Sigma}\rangle$ with g_1 and the $\blacktriangledown_\sigma$ derived from R as above. (It should be clear that g_1 and the $\blacktriangledown_\sigma$ satisfy the requisite properties from Definition 11.) Let's now see an example of $\mathscr{N}_1(R)$ "in action".

Example 18. For this example we again assume our background propositional language contains just two propositional variables, with $\mathscr{W} = \{a,b,c,d\}$. We also assume that $Sources = \{0,1,2\}$. Suppose source 1 gives initial information $\{a\}$, source 2 gives $\{c\}$ and completely reliable source 0 gives $\mathscr{W}$ (and so effectively plays no role in the negotiation). Suppose our ranking assignment R is such that $[R_1(\{a\})]$ and $[R_2(\{c\})]$ are specified as follows (cf. Example 16):

	0	1	2	3
$[R_1(\{a\})]$	a	b	c, d	
$[R_2(\{c\})]$	c		a, d	b

We construct the complete negotiation process $f^{\mathscr{N}_1(R)}(\mathscr{W}, \{a\}, \{c\}) = \sigma$ stage by stage, starting with $\sigma_0 = (\langle \mathscr{W}, \{a\}, \{c\}\rangle)$. Since we have obvious disagreement between sources 1 and 2, a first negotiation round is required. Now we have $[R_1(\{a\})](\overline{\{a\}}) = 1 < 2 = [R_2(\{c\})](\overline{\{c\}})$, i.e., source 1 is less certain of his current standpoint than source 2. Hence we have $g_1(\sigma_0) = \{1\}$, i.e., 1 loses the round and so must weaken. We have $\blacktriangledown_{\sigma_0}(1) = \{a\} \cup \{w \in \overline{\{a\}} \mid [R_1(\{a\})](w)$ is minimal$\}$, i.e., 1 adds to $\{a\}$ the most plausible non-a worlds according to $[R_1(\{a\})]$. Since b is the unique such world, this means $\blacktriangledown_{\sigma_0}(1) = \{a,b\}$ and so we reach the next negotiation stage $\sigma_1 = (\langle \mathscr{W}, \{a\}, \{c\}\rangle, \langle \mathscr{W}, \{a,b\}, \{c\}\rangle)$. Since consistency has still not been reached, another negotiation round is necessary. This time we have $[R_1(\{a\})](\overline{\{a,b\}}) = 2 = [R_2(\{c\})](\overline{\{c\}})$. Hence now both sources are equally certain of their current standpoints. Hence $g_1(\sigma_1) = \{1,2\}$, i.e., both sources must weaken. We have $\blacktriangledown_{\sigma_1}(1) = \{a,b\} \cup \{w \in \overline{\{a,b\}} \mid [R_1(\{a\})](w)$ is minimal$\} = \{a,b,c,d\} = \mathscr{W}$ and $\blacktriangledown_{\sigma_1}(2) = \{c\} \cup \{w \in \overline{\{c\}} \mid [R_2(\{c\})](w)$ is minimal$\} = \{a,c,d\}$. Hence we reach the next stage $\sigma_2 = (\langle \mathscr{W}, \{a\}, \{c\}\rangle, \langle \mathscr{W}, \{a,b\}, \{c\}\rangle, \langle \mathscr{W}, \mathscr{W}, \{a,c,d\}\rangle)$. Since we have now reached consistency, we end the process here with

$$f^{\mathscr{N}_1(R)}(\mathscr{W}, \{a\}, \{c\}) = \sigma_2.$$

From this we deduce $\mathbf{f}^{\mathscr{N}_1(R)}(\mathscr{W}, \{a\}, \{c\}) = \langle \mathscr{W}, \mathscr{W}, \{a, c, d\}\rangle$. For the corresponding merging operator we have

$$\Delta_{\mathscr{N}_1(R)}(\mathscr{W}, \{a\}, \{c\}) = \bigcap_{i=0}^{2} \mathbf{f}_i^{\mathscr{N}_1(R)}(\mathscr{W}, \{a\}, \{c\}) = \{a,c,d\}.$$

As this example illustrates, the combined effect of our g_1 and the $\blacktriangledown_\sigma$ is, roughly speaking, a process in which the sources simultaneously add worlds rank by rank to their initial information until consistency is reached. (See Section A.3 in Appendix A for a precise elaboration of this remark.) In particular, this results in the following behaviour for the individual contraction functions $\ominus_i^{\mathscr{N}_1(R)}$.

Proposition 19. *Let R be a ranking assignment and let $i \in Sources^+$. Then, for all S, $T \in \mathscr{B}$, $S \ominus_i^{\mathscr{N}_1(R)} T = \{w \in \mathscr{W} \mid [R_i(S)](w) \leqslant [R_i(S)](T)\}$.*

In other words, when faced with completely reliable information T, source i weakens his own information S by simply admitting *all* worlds which are at least as plausible as T according to the ranking he provides with S. From this the following can be shown:

Proposition 20. *Let R be a ranking assignment and let $i \in Sources^+$. Then the function $\ominus_i^{\mathscr{N}_1(R)}$ satisfies, in addition to (ind1)–(ind3) from Proposition 4, the following two properties:*

(ind4) $\ominus_i^{\mathscr{N}_1(R)}(T_1 \cup T_2) \subseteq S \ominus_i^{\mathscr{N}_1(R)} T_1$.
(ind5) *If* $(S \ominus_i^{\mathscr{N}_1(R)}(T_1 \cup T_2)) \cap T_1 \neq \emptyset$ *then* $S \ominus_i^{\mathscr{N}_1(R)} T_1 \subseteq S \ominus_i^{\mathscr{N}_1(R)}(T_1 \cup T_2)$.

This means that $\ominus_i^{\mathscr{N}_1(R)}$ belongs to the class of contraction operators known as *severe withdrawal* operators, which were studied in [21]. The rules (ind4) and (ind5) essentially correspond to the postulates ($\dot{-}$7a) and ($\dot{-}$8) given there. In fact (ind5) also corresponds to one of the two AGM *supplementary* contraction postulates [8]. Rule (ind4) is an "antitony" condition, which says strengthening the completely reliable information should result in i having to do more weakening. Taken together, (ind4) and (ind5) say that weakening to accommodate T_1 should produce the same result as weakening to accommodate $T_1 \cup T_2$ provided the weakening which accommodates $T_1 \cup T_2$ already accommodates T_1.

Turning to the merging operator yielded from such a belief negotiation model $\mathcal{N}_1(R)$, we have the following nice characterisation of $\Delta_{\mathcal{N}_1(R)}$.

Proposition 21. *Let R be a ranking assignment. Then, for all* $\vec{S} \in \mathscr{B}^{Sources}$, *we have* $\Delta_{\mathcal{N}_1(R)}(\vec{S}) = \{w \in S_0 \mid \max_{i \in Sources^+}[R_i(S_i)](w) \text{ is minimal}\}$.

This "minimax" operator is a generalised version of the merging operator with integrity constraints Δ^{Max} given in [13], which employs a particular family of ranking assignments based on a notion of (symmetric) *distance* between propositional worlds. Similar operators are also discussed in [17,19,20], and are shown to satisfy several interesting properties.

How do the SC functions $\mathbf{f}^{\mathcal{N}_1(R)}$ fare with regard to the minimal change postulates from Section 2.1? Well quite badly as it turns out. Indeed they do not, in general, satisfy even either of the weaker postulates (sc6) and (sc7) mentioned there. The ranking assignment R used in Example 18 provides a counter-example against (sc6). To see this note that, in that example, we have

$$\{a\} \cap \mathbf{f}_0^{\mathcal{N}_1(R)}(\mathscr{W}, \{a\}, \{c\}) \cap \mathbf{f}_2^{\mathcal{N}_1(R)}(\mathscr{W}, \{a\}, \{c\}) = \{a\} \cap \mathscr{W} \cap \{a, c, d\} \neq \emptyset.$$

Now if $\mathbf{f}^{\mathcal{N}_1(R)}$ satisfied (sc6) we would conclude

$$\mathbf{f}_1^{\mathcal{N}_1(R)}(\mathscr{W}, \{a\}, \{c\}) \subseteq \{a\} \ominus_1^{\mathcal{N}_1(R)} \overline{\{a\}}.$$

But $\mathbf{f}_1^{\mathcal{N}_1(R)}(\mathscr{W}, \{a\}, \{c\}) = \mathscr{W}$ and $\{a\} \ominus_1^{\mathcal{N}_1(R)} \overline{\{a\}} = \{a, b\}$. Hence $\mathbf{f}^{\mathcal{N}_1(R)}$ does not satisfy (sc6). That the $\mathbf{f}^{\mathcal{N}_1(R)}$ don't validate (sc7) can be shown by the following counter-example.

Example 22. As in Example 18 we again assume $\mathscr{W} = \{a, b, c, d\}$, *Sources* = {0, 1, 2}, and that *Sources* provide the information profile $\vec{S} = (\mathscr{W}, \{a\}, \{c\})$. This time, however, let the ranking assignment R be such that $[R_1(\{a\})]$ and $[R_2(\{c\})]$ are specified as follows

	0	1	2
$[R_1(\{a\})]$	a	c	b, d
$[R_2(\{c\})]$	c	b	a, d

Then it can be checked that $\mathbf{f}^{\mathcal{N}_1(R)}(\vec{S}) = (\mathscr{W}, \{a, c\}, \{b, c\})$. Clearly $\mathbf{f}_2^{\mathcal{N}_1(R)}(\vec{S}) \neq S_2$, hence if $\mathbf{f}^{\mathcal{N}_1(R)}$ satisfied (sc7) we would deduce that there is some consistent $\vec{T}$ such that $\vec{S} \subseteq \vec{T} \subseteq \mathbf{f}^{\mathcal{N}_1(R)}(\vec{S})$ and $S_2 \cap T_0 \cap T_1 = \emptyset$. Since $\vec{S} \subseteq \vec{T}$ we must have $T_0 = \mathscr{W}$ and so, since $S_2 = \{c\}$, this latter amounts to saying $c \notin T_1$. But it is straightforward to see that if $\vec{T} \subseteq \mathbf{f}^{\mathcal{N}_1(R)}(\vec{S})$ and $c \notin T_1$ then $\vec{T}$ must be inconsistent. Hence (sc7) cannot hold.

It would be interesting to find out if there are any additional conditions we could place on g_2 or on the $\blacktriangledown_\sigma$ which could help to capture (sc7) for $\mathbf{f}^{\mathcal{N}_1(R)}$.

Since, as we remarked at the end of Section 2.1, the "equilibrium" property (sc8) implies (sc5) (and therefore also (sc6) and (sc7)), this means that (sc8) also fails to hold for $\mathbf{f}^{\mathcal{N}_1(R)}$. However, we can at least show that the $\mathbf{f}^{\mathcal{N}_1(R)}$ do satisfy "one half" of (sc8).

Proposition 23. *Let R be a ranking assignment. Then, for all* $i \in Sources^+$ *and all* $\vec{S} \in \mathscr{B}^{Sources}$, *we have* $\mathbf{f}_i^{\mathcal{N}_1(R)}(\vec{S}) \supseteq S_i \ominus_i^{\mathcal{N}_1(R)} \left(\bigcap_{j \neq i} \mathbf{f}_j^{\mathcal{N}_1(R)}(\vec{S})\right)$.

Summing up, it seems, interestingly, that, while $\Delta_{\mathcal{N}_1(R)}$ might be quite well-behaved, there still seems to be room for improvement regarding the behaviour of $\mathbf{f}^{\mathcal{N}_1(R)}$.

4.2. Second instantiation

Our second instantiation of the framework is about taking a more *orderly* approach to the negotiation process. The idea now is that the sources in *Sources*$^+$ each take it in turn to weaken their information according to some given fixed running order. Each source, during his turn, repeatedly weakens his information until it becomes jointly consistent with the information of all the sources who have taken their turn already. This amounts to fixing $\mathbf{f}^{\mathcal{N}}(\vec{S})$ one element at a time, starting with $\mathbf{f}_0^{\mathcal{N}}(\vec{S}) = S_0$. So, using $\prec$ to denote a given strict total order on *Sources*$^+$ and assuming $i_1 \prec i_2 \prec \cdots \prec i_n$, we first focus on i_1 and repeatedly weaken S_{i_1} until it becomes consistent with S_0. The result of this weakening we will take to be $\mathbf{f}_{i_1}^{\mathcal{N}}(\vec{S})$. Of course it may be that $S_{i_1} \cap S_0 \neq \emptyset$ to begin with, in which case i_1 needn't do any weakening at all. Next we focus on i_2 and repeatedly weaken S_{i_2} until it becomes consistent with $\mathbf{f}_{i_1}^{\mathcal{N}}(\vec{S}) \cap S_0$. The result of *this* weakening we will take to be $\mathbf{f}_{i_2}^{\mathcal{N}}(\vec{S})$. Then it is the turn of i_3, and so on through the rest of the sources. For simplicity, and without loss of generality, in what follows we shall take $\prec$ to be just the usual ordering $<$ on the natural numbers, i.e., we assume source 1 weakens first, followed by source 2, then source 3, and so on.

To fit this idea into our framework we need to define suitable functions g and $\blacktriangledown_\sigma$. For the former we define the function $g_2 : \Sigma \to 2^{Sources^+}$ by setting, for each negotiation stage $\sigma = (\vec{S}^0, \ldots, \vec{S}^m)$,

$g_2(\sigma) = \{i\}$, where
$i \in Sources^+$ is minimal such that $\bigcap_{j \leqslant i} S_j^m = \emptyset$.

For the $\blacktriangledown_\sigma$ we shall assume the weakenings are carried out in exactly the same manner as before with the help of a given ranking assignment R. Thus we define the belief negotiation model $\mathcal{N}_2(R) = \langle g_2, \{\blacktriangledown_\sigma\}_{\sigma \in \Sigma} \rangle$ where now g_2 is defined as above and the $\blacktriangledown_\sigma$ are defined from R as in the previous subsection. (Again it is obvious that g_2 satisfies the requisite properties from Definition 11.) Let us give a worked example of a belief negotiation model of this type.

Example 24. Suppose once more that $\mathcal{W} = \{a, b, c, d\}$, but this time that $Sources = \{0, 1, 2, 3\}$. We suppose that our sources provide the information profile $\vec{S} = (\{a, b, c\}, \{d\}, \{a, b, d\}, \{c\})$. We will use the belief negotiation model $\mathcal{N}_2(R)$, where R is such that $[R_1(\{d\})]$, $[R_2(\{a, b, d\})]$ and $[R_3(\{c\})]$ are given as follows:

	0	1	2	3
$[R_1(\{d\})]$	d	a, b	c	
$[R_2(\{a, b, d\})]$	a, b, d		c	
$[R_3(\{c\})]$	c	d	a	b

Let us construct the sequence $f^{\mathcal{N}_2(R)}(\vec{S}) = \sigma$ stage by stage, starting with $\sigma_0 = (\vec{S}^0)$ where $\vec{S}^0 = \vec{S}$. Clearly $\vec{S}^0$ is inconsistent, so a first negotiation round is necessary. According to the definition of g_2, determining who must weaken at this initial negotiation stage is a matter of going through each of the sources in $Sources^+$ in the order prescribed by < and selecting the first one for which $\bigcap_{j \leqslant i} S_j^0 = \emptyset$. Starting then with source 1, we immediately see that $\bigcap_{j \leqslant 1} S_j^0 = S_0^0 \cap S_1^0 = \{a, b, c\} \cap \{d\} = \emptyset$. Hence source 1 is the loser of this negotiation round, i.e., $g_2(\sigma_0) = \{1\}$, and so must make some weakening. Since $\blacktriangledown_{\sigma_0}(1) = \{d\} \cup \{w \in \overline{\{d\}} \mid [R_1(\{d\})](w)$ is minimal$\} = \{a, b, d\}$ this leads us to the next stage $\sigma_1 = (\vec{S}^0, \vec{S}^1)$, where $\vec{S}^1 = (\{a, b, c\}, \{a, b, d\}, \{a, b, d\}, \{c\})$. Since consistency has not yet been reached, a second negotiation round is necessary. As a result of his weakening at the previous stage, source 1's current standpoint is no longer in conflict with that of source 0, i.e., we have $S_0^1 \cap S_1^1 = \{a, b, c\} \cap \{a, b, d\} \neq \emptyset$. Hence source 1 weakens no further. We must consider source 2 next. But $\bigcap_{j \leqslant 2} S_j^2 = S_0^1 \cap S_1^1 \cap S_2^1 = \{a, b, c\} \cap \{a, b, d\} \cap \{a, b, d\} \neq \emptyset$ and so 2 need not weaken either. Since source 3 is the only source left, this means we must have $g_2(\sigma_1) = \{3\}$. Now $\blacktriangledown_{\sigma_1}(3) = \{c\} \cup \{w \in \overline{\{c\}} \mid [R_3(\{c\})](w)$ is minimal$\} = \{c, d\}$ which leads us to the next stage $\sigma_2 = (\vec{S}^0, \vec{S}^1, \vec{S}^2)$ where $\vec{S}^2 = (\{a, b, c\}, \{a, b, d\}, \{a, b, d\}, \{c, d\})$. Since we have still not reached consistency, source 3 is required to do yet more weakening, i.e., we have $g_2(\sigma_2) = \{3\}$. This time we have $\blacktriangledown_{\sigma_2}(3) = \{c, d\} \cup \{w \in \overline{\{c, d\}} \mid [R_3(\{c\})](w)$ is minimal$\} = \{a, c, d\}$ leading to the next stage $\sigma_3 = (\vec{S}^0, \vec{S}^1, \vec{S}^2, \vec{S}^3)$ where now $\vec{S}^3 = (\{a, b, c\}, \{a, b, d\}, \{a, b, d\}, \{a, c, d\})$. This time we *have* reached consistency, so the process stops here with $f^{\mathcal{N}_2(R)}(\vec{S}) = \sigma_3$ and $\mathbf{f}^{\mathcal{N}_2(R)}(\vec{S}) = \vec{S}^3 = (\{a, b, c\}, \{a, b, d\}, \{a, b, d\}, \{a, c, d\})$. For the corresponding merging operator we get $\Delta_{\mathcal{N}_2(R)}(\vec{S}) = \bigcap_{i=0}^{3} S_i^3 = \{a\}$.

Note that, by Proposition 14, the $\ominus_i^{\mathcal{N}_2(R)}$ are the same as the $\ominus_i^{\mathcal{N}_1(R)}$ from the previous subsection. What can we say this time about the SC functions $\mathbf{f}^{\mathcal{N}_2(R)}$? First of all we may show the following.

Proposition 25. *Let R be a ranking assignment and let $i \in Sources^+$. Then, for each information profile $\vec{S}$, we have*

$$\mathbf{f}_i^{\mathcal{N}_2(R)}(\vec{S}) = S_i \ominus_i^{\mathcal{N}_2(R)} \left(\bigcap_{j<i} \mathbf{f}_j^{\mathcal{N}_2(R)}(\vec{S}) \right).$$

In other words $\mathbf{f}_i^{\mathcal{N}_2(R)}(\vec{S})$ is equal to the result—according to i's individual contraction function relative to $\mathbf{f}^{\mathcal{N}_2(R)}$—of weakening S_i to be jointly consistent with all the $\mathbf{f}_j^{\mathcal{N}_2(R)}(\vec{S})$ for which $j < i$. Using this together with the fact that the $\ominus_i^{\mathcal{N}_2(R)}$ satisfy the properties (ind4) and (ind5) from Proposition 20 then allows us to prove:

Proposition 26. *Let R be a ranking assignment. Then the SC function $\mathbf{f}^{\mathcal{N}_2(R)}$ satisfies (sc8).*

Thus, imposing a strict "order of weakening" on the sources has forced our SC function to satisfy the equilibrium property (sc8) (and hence also (sc5), (sc6) and (sc7)). Meanwhile we can characterise $\Delta_{\mathcal{N}_2(R)}$ with the help of the following piece of extra notation: We let $<_{lex}$ denote the lexicographic ordering on $\mathbb{N}^n$, i.e., given two tuples $\vec{x}, \vec{y} \in \mathbb{N}^n$ such that $\vec{x} = (x_1, \ldots, x_n)$ and $\vec{y} = (y_1, \ldots, y_n)$, we have $\vec{x} <_{lex} \vec{y}$ iff there exists j such that (i) $x_j < y_j$ and (ii) $x_i = y_i$ for all $i < j$. (Clearly $<_{lex}$ is a strict total order on $\mathbb{N}^n$.) Then we have the following.

Proposition 27. *Let R be a ranking assignment. Then, using r_j as an abbreviation for $[R_j(S_j)]$, we have*

$$\Delta_{\mathcal{N}_2(R)}(\vec{S}) = \{w \in S_0 \mid (r_1(w), r_2(w), \ldots, r_n(w)) \text{ is minimal under } <_{lex}\}.$$

Thus $\Delta_{\mathcal{N}_2(R)}(\vec{S})$ collects all the "best" worlds in S_0, in the special sense where one world is considered "better" than another if it is assigned lower rank by source 1, or, in case they are assigned the same rank by 1, it is assigned a lower rank by 2, or, in case they are also assigned the same rank by 2, it is assigned a lower rank by 3, or, etc. Thus the effect when merging is that the opinion of source i is given *precedence* over that of i' whenever $i < i'$. Such a lexicographic approach to merging has been considered in [17] (see Section 4.5 there) where the ordering < on the sources is interpreted as a given ordering of *reliability* on the sources, i.e., the most reliable sources are given precedence.

5. Conclusion

We have made a start on the study of social contraction functions, which are applicable to the problem of merging information from multiple sources. The intention is that social contraction is to merging what contraction is to belief revision. We have considered both a postulational and a procedural approach, managing in the process of the latter to extend the belief negotiation framework of [5]. Our investigations are at an early stage, and much still needs to be done. From the postulational viewpoint we feel there are still many more postulates for social contraction waiting to be discovered and evaluated. From the negotiation viewpoint we looked in this paper at only two relatively simple possible ways of instantiating the basic negotiation framework. We are presently looking at various other, more complex, ways in which this can be done. One suggestion, due to Thomas Meyer, relates to the $\blacktriangledown_\sigma$-functions. Instead of blindly adding *all* the most plausible worlds not yet in source i's current standpoint S_i^m as is done in this paper, the function $\blacktriangledown_\sigma(i)$ should be more selective and add only those which are already included in at least one of the current standpoints S_j^m of the other sources at stage σ. (If none of these most plausible worlds appear in *any* of the S_j^m then $\blacktriangledown_\sigma(i)$ should add all of them as before.) Refinements such as this could lead to more interesting social contraction behaviour. Finally, we would also like to explore more fully the relationship between the merging operators derived from social contraction and the integrity constraints merging operators of [13]. In particular, it would be interesting to find out whether any of the additional minimal change SC postulates from Section 2.1 induce corresponding postulates for the derived merging operators.

Acknowledgements

This work is supported by the DFG research project "Computationale Dialektik". Thanks are due to Thomas Meyer for a couple of helpful suggestions, and to Alexander Nittka and three anonymous referees of this journal, whose comments helped to improve the paper.

Appendix A. Proofs

A.1. Proofs from Section 2

Proof of Proposition 4. Let $S, T \in \mathscr{B}$ and let $\vec{U}$ be the information profile such that $U_i = S$, $U_0 = T$ and $U_j = \mathscr{W}$ for all $j \notin \{0, i\}$. Then, we have $\mathbf{f}_i(\vec{S}) = S \ominus_i^{\mathbf{f}} T$ (by definition of $\ominus_i^{\mathbf{f}}$), $\mathbf{f}_0(\vec{U}) = T$ (by (sc4)) and $\mathbf{f}_j(\vec{U}) = \mathscr{W}$ for all $j \notin \{0, i\}$ (by (sc1)). Then to show (ind1) we have $U_i \subseteq \mathbf{f}_i(\vec{U})$ by (sc1), i.e., $S \subseteq S \ominus_i^{\mathbf{f}} T$ as required. For (ind2) we know $\mathbf{f}(\vec{U})$ is consistent by (sc2), i.e., $\bigcap_k \mathbf{f}_k(\vec{U}) \neq \emptyset$. But $\bigcap_k \mathbf{f}_k(\vec{U}) = (S \ominus_i^{\mathbf{f}} T) \cap T$, which gives the required conclusion. Finally for (ind3) suppose $S \cap T = \cap_k U_k \neq \emptyset$. Using (sc3) we deduce that $\mathbf{f}(\vec{U}) = \vec{U}$, in particular $U_i = \mathbf{f}_i(\vec{U})$, i.e., $S = S \ominus_i^{\mathbf{f}} T$ as required. □

Proof of Proposition 5. Suppose $\mathbf{f}$ satisfies (sc1), (sc4) and (sc5). To show (sc3), suppose $\vec{S}$ is consistent. We must show $\mathbf{f}_i(\vec{S}) = S_i$ for all $i \in Sources$. If $i = 0$ then this holds from (sc4). So suppose $i \in Sources^+$. Since $\vec{S}$ is consistent we know $S_i \cap \cap_{j \neq i} S_j \neq \emptyset$. But, using (sc1), we have $\bigcap_{j \neq i} S_j \subseteq \bigcap_{j \neq i} \mathbf{f}_j(\vec{S})$. Hence, from $S_i \cap \cap_{j \neq i} S_j \neq \emptyset$ we may deduce $S_i \cap \bigcap_{j \neq i} \mathbf{f}_j(\vec{S}) \neq \emptyset$. Applying (sc5) to this then gives $\mathbf{f}_i(\vec{S}) = S_i$ as required. □

Proof of Proposition 7. For the first part, suppose $\mathbf{f}$ satisfies (sc4) and (sc5+). To show (sc5), let $i \in Sources^+$ and suppose $\mathbf{f}_i(\vec{S}) \neq S_i$. We must show $S_i \cap \bigcap_{j \neq i} \mathbf{f}_j(\vec{S}) = \emptyset$. But from $\mathbf{f}_i(\vec{S}) \neq S_i$ we know $i \notin X_{\mathbf{f}}(\vec{S})$. This tells us $\{j \in Sources \mid j \neq i\} \supseteq X_{\mathbf{f}}(\vec{S}) \cup \{0\}$ and so $S_i \cap \bigcap_{j \neq i} \mathbf{f}_j(\vec{S}) \subseteq S_i \cap (\bigcap_{j \in X_{\mathbf{f}}(\vec{S})} \mathbf{f}_j(\vec{S})) \cap \mathbf{f}_0(\vec{S})$. By definition of $X_{\mathbf{f}}(\vec{S})$ we know $\mathbf{f}_j(\vec{S}) = S_j$ for all $j \in X_{\mathbf{f}}(\vec{S})$, while also $\mathbf{f}_0(\vec{S}) = S_0$ by (sc4). Hence $S_i \cap \bigcap_{j \neq i} \mathbf{f}_j(\vec{S}) \subseteq S_i \cap (\bigcap_{j \in X_{\mathbf{f}}(\vec{S})} S_j) \cap S_0$. From $i \notin X_{\mathbf{f}}(\vec{S})$ we have $S_i \cap (\bigcap_{j \in X_{\mathbf{f}}(\vec{S})} S_j) \cap S_0 = \emptyset$ by (sc5+). Hence $S_i \cap \bigcap_{j \neq i} \mathbf{f}_j(\vec{S}) = \emptyset$ as required.

For the second part, suppose $\mathbf{f}$ satisfies (sc4), (sc5) and (sc1+). Let $i \in Sources$. Then, using (**sc**1+) and (sc4) allows us to write $\bigcap_{j \neq i} \mathbf{f}_j(\vec{S}) = (\bigcap_{j \in X_{\mathbf{f}}(\vec{S})} S_j) \cap S_0$. Hence we see the antecedents of (sc5) and (sc5+) are equivalent. Since the consequents of the two rules are clearly also equivalent, the result follows. □

Proof of Proposition 8. To show (sc3) suppose $\vec{S}$ is consistent. We must show $\mathbf{f}_i(\vec{S}) = S_i$ for all $i \in Sources$. The case $i = 0$ is handled by (sc4), so let $i \in Sources^+$. Suppose for contradiction that $\mathbf{f}_i(\vec{S}) \neq S_i$. Then (sc7) tells us there exists some consistent $\vec{T}$ such that $\vec{S} \subseteq \vec{T} \subseteq \mathbf{f}(\vec{S})$ and $S_i \cap \bigcap_{j \neq i} T_j = \emptyset$. But from $\vec{S} \subseteq \vec{T}$ we get $\bigcap_{j \neq i} S_j \subseteq \bigcap_{j \neq i} T_j$, and so from $S_i \cap \bigcap_{j \neq i} T_j = \emptyset$ we get $S_i \cap \bigcap_{j \neq i} S_j = \emptyset$, i.e., $\vec{S}$ is inconsistent—contradiction. Hence $\mathbf{f}_i(\vec{S}) = S_i$ as required. □

Proof of Proposition 10. Let $\mathbf{f}$ be an SC function. Let $i \in Sources^+$ and suppose $S_i \cap \bigcap_{j \neq i} \mathbf{f}_j(\vec{S}) \neq \emptyset$. To show (sc5) we must show $\mathbf{f}_i(\vec{S}) = S_i$. But since $\ominus_i^{\mathbf{f}}$ satisfies (ind3) by Proposition 4 we have $S_i \ominus_i^{\mathbf{f}}(\bigcap_{j\neq i}\mathbf{f}_j(\vec{S})) = S_i$. Hence the assumption that $\mathbf{f}_i(\vec{S}) \subseteq S_i \ominus_i^{\mathbf{f}}(\bigcap_{j\neq i}\mathbf{f}_j(\vec{S}))$ yields $\mathbf{f}_i(\vec{S}) \subseteq S_i$. We obtain equality by (sc1). □

A.2. Proofs from Section 3

Proof of Theorem 13. First we show that $\mathbf{f}^{\mathcal{N}}$ is an SC function for any belief negotiation model $\mathcal{N}$. We need to show that $\mathbf{f}^{\mathcal{N}}$ satisfies (sc1)–(sc4). For (sc1) let $\vec{S} \in \mathscr{B}^{Sources}$ and suppose $f^{\mathcal{N}}(\vec{S}) = (\vec{S}^0, \ldots, \vec{S}^k)$. We must show $\vec{S} \subseteq \mathbf{f}^{\mathcal{N}}(\vec{S})$, i.e., $\vec{S}^0 \subseteq \vec{S}^k$. But, for each $0 \leqslant j < k$, we have $\vec{S}^j \subseteq \vec{S}^{j+1}$ (this is ensured by (▼0a)) and so, since the "⊆" relation between information profiles is clearly transitive, we get the required conclusion. Since we always have $\vec{S}^k$ is consistent this means $\mathbf{f}^{\mathcal{N}}(\vec{S})$ is consistent and so (sc2) also holds. For (sc3) we have that if $\vec{S}$ is consistent then we must have $f^{\mathcal{N}}(\vec{S}) = (\vec{S})$ and so $\mathbf{f}^{\mathcal{N}}(\vec{S}) = \vec{S}$ as required. Finally for (sc4), since $0 \notin g(\sigma)$ for all σ, we clearly have $S_0^j = S_0^0 = S_0$ for all $0 \leqslant j \leqslant k$. In particular we have $S_0^k = S_0$, i.e., $\mathbf{f}_0(\vec{S}) = S_0$ as required. Hence $\mathbf{f}^{\mathcal{N}}$ is indeed an SC function.

Now we show that, given an SC function $\mathbf{f}$, there exists a belief negotiation model $\mathcal{N} = \langle g, \{\blacktriangledown_\sigma\}_{\sigma\in\Sigma}\rangle$ such that $\mathbf{f} = \mathbf{f}^{\mathcal{N}}$. We define the functions g and $\blacktriangledown_\sigma$ from $\mathbf{f}$ in turn and then show that $\mathbf{f} = \mathbf{f}^{\mathcal{N}}$.

Defining g. Given $\mathbf{f}$, we define the function $g : \Sigma \to 2^{Sources^+}$ by setting, for each $\sigma = (\vec{S}^0, \ldots, \vec{S}^m) \in \Sigma$,

$$g(\sigma) = \{i \in Sources^+ \mid S_i^m \neq \mathbf{f}_i(\vec{S}^m)\}.$$

We need to check that g so defined satisfies the conditions (g0a) and (g0b). To show (g0a) is satisfied, i.e., that $g(\sigma) \neq \emptyset$, note first that $\vec{S}^m$ is inconsistent by definition of the set Σ. Now suppose for contradiction that $g(\sigma) = \emptyset$. Then we must have $S_i^m = \mathbf{f}_i(\vec{S}^m)$ for all $i \in Sources^+$. Since we additionally have $S_0^m = \mathbf{f}_0(\vec{S}^m)$ by (sc4), this means $S_i^m = \mathbf{f}_i(\vec{S}^m)$ for all $i \in Sources$, i.e., $\vec{S}^m = \mathbf{f}(\vec{S}^m)$. Hence, since $\mathbf{f}(\vec{S}^m)$ is consistent by (sc2), this gives us that $\vec{S}^m$ is consistent—contradiction. Hence $g(\sigma) \neq \emptyset$ as required. Turning to (g0b), we must show that $i \in g(\sigma)$ implies $S_i^m \neq \mathscr{W}$. But if $S_i^m = \mathscr{W}$ then, by (sc1), we must have $S_i^m = \mathbf{f}_i(\vec{S}^m)$ and so $i \notin g(\sigma)$ as required.

Defining the $\blacktriangledown_\sigma$. For each $\sigma = (\vec{S}^0, \ldots, \vec{S}^m) \in \Sigma$ we define the function $\blacktriangledown_\sigma : Sources^+ \to \mathscr{B}$ by setting, for each $i \in Sources^+$,

$$\blacktriangledown_\sigma(i) = \begin{cases} \mathbf{f}_i(\vec{S}^m) & \text{if } S_i^m \neq \mathbf{f}_i(\vec{S}^m) \\ \mathscr{W} & \text{otherwise.} \end{cases}$$

We now need to check that the $\blacktriangledown_\sigma$ so defined satisfy the properties (▼0a) and (▼0b). That (▼0a) is satisfied, i.e., $S_i^m \subseteq \blacktriangledown_\sigma(i)$, follows almost immediately from (sc1). For (▼0b) we must show that $\blacktriangledown_\sigma(i) = S_i^m$ implies $S_i^m = \mathscr{W}$. So suppose $\blacktriangledown_\sigma(i) = S_i^m$. Then obviously it cannot be the case that both $\blacktriangledown_\sigma(i) = \mathbf{f}_i(\vec{S}^m)$ and $S_i^m \neq \mathbf{f}_i(\vec{S}^m)$. This rules out the first clause in the definition of $\blacktriangledown_\sigma(i)$ and so it must be that we are in the second clause, i.e., that $\blacktriangledown_\sigma(i) = \mathscr{W}$ (and $S_i^m = \mathbf{f}_i(\vec{S}^m)$). Hence, since we assumed $\blacktriangledown_\sigma(i) = S_i^m$, we have $S_i^m = \mathscr{W}$ as required.

Given $\mathcal{N}$ defined above, it remains to show that $\mathbf{f}(\vec{S}) = \mathbf{f}^{\mathcal{N}}(\vec{S})$ for all information profiles $\vec{S}$. We will show this by first constructing, for a given $\vec{S}$, the sequence $f^{\mathcal{N}}(\vec{S})$ representing the complete process of negotiation on $\vec{S}$. For the case when $\vec{S}$ is consistent we clearly have $f^{\mathcal{N}}(\vec{S}) = (\vec{S})$ and so $\mathbf{f}^{\mathcal{N}}(\vec{S}) = \vec{S}$. Since in this case we know also $\mathbf{f}(\vec{S}) = \vec{S}$ by (sc3) we get $\mathbf{f}(\vec{S}) = \mathbf{f}^{\mathcal{N}}(\vec{S})$ as required. So suppose now that $\vec{S}$ is inconsistent. In this case, we claim that $f^{\mathcal{N}}(\vec{S}) = (\vec{S}, \mathbf{f}(\vec{S}))$. To see this, let $\sigma_0 = (\vec{S}) \in \Sigma$ denote the initial negotiation stage and let $\sigma_1 = (\vec{S}, \vec{S}^1)$ denote the stage which follows the first negotiation round. We show that $\vec{S}^1 = \mathbf{f}(\vec{S})$. First, since $0 \notin g(\sigma_0)$ as always, we have $S_0^1 = S_0 =$ (by (sc4)) $\mathbf{f}_0(\vec{S})$. So now let $i \in Sources^+$. If $i \in g(\sigma_0)$ then i must weaken and so $S_i^1 = \blacktriangledown_{\sigma_0}(i)$. By definition of g, we have $S_i \neq \mathbf{f}_i(\vec{S})$ and so, by definition of $\blacktriangledown_{\sigma_0}$, we have $S_i^1 = \mathbf{f}_i(\vec{S})$. If $i \notin g(\sigma_0)$ then i does not weaken, i.e., $S_i^1 = S_i$. By definition of g we have $S_i = \mathbf{f}_i(\vec{S})$. Hence again $S_i^1 = \mathbf{f}_i(\vec{S})$ as required. Hence $\vec{S}^1 = \mathbf{f}(\vec{S})$. Since $\mathbf{f}(\vec{S})$ is consistent by (sc2), the negotiation process ends here with $f^{\mathcal{N}}(\vec{S}) = (\vec{S}, \mathbf{f}(\vec{S}))$ and so $\mathbf{f}^{\mathcal{N}}(\vec{S}) = \mathbf{f}(\vec{S})$ as required. □

Proof of Proposition 14. We must show that $S \ominus_i^{\mathcal{N}} T = S \ominus_i^{\mathcal{N}'} T$ for all $S, T \in \mathscr{B}^{Sources}$. So, given S and T, let $\vec{U} \in \mathscr{B}^{Sources}$ be such that $U_i = S$, $U_0 = T$ and $U_j = \mathscr{W}$ for $j \notin \{0, i\}$. Then we have $S \ominus_i^{\mathcal{N}} T = \mathbf{f}_i^{\mathcal{N}}(\vec{U})$ and $S \ominus_i^{\mathcal{N}'} T = \mathbf{f}_i^{\mathcal{N}'}(\vec{U})$. Hence we must show $\mathbf{f}_i^{\mathcal{N}}(\vec{U}) = \mathbf{f}_i^{\mathcal{N}'}(\vec{U})$. We will show that in fact $f^{\mathcal{N}}(\vec{U}) = f^{\mathcal{N}'}(\vec{U})$, which clearly suffices. So let $\sigma^{\mathcal{N}} = f^{\mathcal{N}}(\vec{U}) = (\vec{U}^0, \ldots, \vec{U}^k)$ and $\sigma^{\mathcal{N}'} = f^{\mathcal{N}'}(\vec{U}) = (\vec{V}^0, \ldots, \vec{V}^l)$. We will first prove by induction on m that $\sigma_m^{\mathcal{N}} = \sigma_m^{\mathcal{N}'}$ for all $0 \leqslant m \leqslant \min\{k, l\}$. For the case $m = 0$ we have $\vec{U}^0 = \vec{U} = \vec{V}^0$ and so $\sigma_0^{\mathcal{N}} = (\vec{U}^0) = (\vec{V}^0) = \sigma_m^{\mathcal{N}'}$ as required. Now suppose $0 < m \leqslant \min\{k, l\}$ and that $\sigma_{m-1}^{\mathcal{N}} = \sigma_{m-1}^{\mathcal{N}'}$, i.e., that $\vec{U}^s = \vec{V}^s$ for all $s \leqslant m-1$. We must show $\sigma_m^{\mathcal{N}} = \sigma_m^{\mathcal{N}'}$, i.e., that additionally $\vec{U}^m = \vec{V}^m$. Since $\vec{U} \subseteq \vec{U}^{m-1} = \vec{V}^{m-1}$ we know $U_j^{m-1} = V_j^{m-1} = \mathscr{W}$ for$j \notin \{0, i\}$. Hence, since g and g' satisfy (g0b), this means that at stages $\sigma_{m-1}^{\mathcal{N}}$ and $\sigma_{m-1}^{\mathcal{N}'}$, respectively, neither g nor g' selects any source $j \neq i$. Since g and g' satisfy (g0a), this means we must have $g(\sigma_{m-1}^{\mathcal{N}}) = g'(\sigma_{m-1}^{\mathcal{N}'}) = \{i\}$. Hence, for $j \notin \{0, i\}$ we have again $U_j^m = V_j^m = \mathscr{W}$, while $U_i^m = \blacktriangledown_{\sigma_{m-1}^{\mathcal{N}}}(i) = \blacktriangledown_{\sigma_{m-1}^{\mathcal{N}'}}(i) = V_i^m$. Meanwhile $U_0^m = V_0^m = T$, hence we have that $U_j^m = V_j^m$ for *all* $j \in Sources$, i.e., $\vec{U}^m = \vec{V}^m$ as required. This completes the inductive step, and so we have shown that $\sigma_m^{\mathcal{N}} = \sigma_m^{\mathcal{N}'}$ for all $0 \leqslant m \leqslant \min\{k, l\}$. Since k, respectively, l, are minimal such that $\vec{U}^k$, respectively $\vec{V}^l$ are consistent, we must have $k = l$. Hence we have $\sigma^{\mathcal{N}} = \sigma^{\mathcal{N}'}$, i.e., $f^{\mathcal{N}}(\vec{U}) = f^{\mathcal{N}'}(\vec{U})$. This completes the proof. □

A.3. Proofs from Section 4.1

For the remaining proofs in the paper it is useful first to introduce some extra notation. For each ranking assignment R and each $x \in \mathbb{N}$ we define a function $\mathbf{h}^{R,x} : \mathscr{B}^{Sources} \to \mathscr{B}^{Sources}$ by setting, for each $\vec{S} \in \mathscr{B}^{Sources}$,

$$\mathbf{h}_i^{R,x}(\vec{S}) = \begin{cases} \{w \in \mathscr{W} \mid [R_i(S_i)](w) \leqslant x\} & \text{if } i \neq 0 \\ S_i & \text{if } i = 0. \end{cases}$$

In other words we have $\mathbf{h}_0^{R,x}(\vec{S}) = S_0$ while, for each $i \in Sources^+$, the entry $\mathbf{h}_i^{R,x}(\vec{S})$ collects all those worlds which are awarded a ranking of at most x by the ranking i provides with S_i. Note that the fact that each ranking $[R_i(S_i)]$ is anchored on S_i means that we have $\mathbf{h}^{R,0}(\vec{S}) = \vec{S}$. For the belief negotiation model $\mathscr{N}_1(R)$ we can neatly describe the negotiation process in terms of these functions. First we give the following lemma, which describes the transition from one negotiation stage to the next.

Lemma 28. *Let R be a ranking assignment and $\vec{S}$ be an information profile. Suppose $f^{\mathscr{N}_1(R)}(\vec{S}) = \sigma = (\vec{S}^0, \ldots, \vec{S}^k)$. Then, for each $m = 0, 1, \ldots, k-1$, if $\vec{S}^m = \mathbf{h}^{R,x}(\vec{S})$ for some x then $\vec{S}^{m+1} = \mathbf{h}^{R,y}(\vec{S})$ where y is minimal such that $\vec{S}^m \subset \mathbf{h}^{R,y}(\vec{S})$.*

Proof. Assume m is such that $\vec{S}^m = \mathbf{h}^{R,x}(\vec{S})$ for some x. For this proof let us abbreviate i's given ranking $[R_i(S_i)]$ by just r_i for each $i \in Sources^+$. Then let $y' = \min_{i \in Sources^+}\{r_i(\overline{S_i^m})\}$. First we claim $\vec{S}^{m+1} = \mathbf{h}^{R,y'}(\vec{S})$. To show this we need to show that $S_i^{m+1} = \mathbf{h}_i^{R,y'}(\vec{S})$ for all $i \in Sources^+$ (clearly we already have $S_0^{m+1} = S_0 = \mathbf{h}_0^{R,y'}(\vec{S})$). There are two cases we need to check: (a) $i \in g_1(\sigma_m)$ and (b) $i \notin g_1(\sigma_m)$. If $i \in g_1(\sigma_m)$, equivalently (by definition of g_1) $r_i(\overline{S_i^m}) = y'$, then we have $S_i^{m+1} = \blacktriangledown_{\sigma_m}(i) = S_i^m \cup \{w \in \overline{S_i^m} \mid r_i(w) \text{ is minimal}\} = S_i^m \cup \{w \in \overline{S_i^m} \mid r_i(w) = y'\}$. By the minimality of y' we know there is no $w \in \overline{S_i^m}$ such that $r_i(w) < y'$. Hence we may just as well write

$$S_i^{m+1} = S_i^m \cup \{w \in \overline{S_i^m} \mid r_i(w) \leqslant y'\}.$$

Since $S_i^m = \mathbf{h}_i^{R,x}(\vec{S})$ we may re-write this as

$$\begin{aligned} S_i^{m+1} &= \{w \in \mathscr{W} \mid r_i(w) \leqslant x\} \cup \{w \in \mathscr{W} \mid x < r_i(w) \leqslant y'\} \\ &= \{w \in \mathscr{W} \mid r_i(w) \leqslant y'\} = \mathbf{h}_i^{R,y'}(\vec{S}) \quad \text{as required.} \end{aligned}$$

For the case $i \notin g(\sigma_m)$, equivalently $r_i(\overline{S_i^m}) > y'$, we have $S_i^{m+1} = S_i^m$. Using $r_i(\overline{S_i^m}) > y'$ together with the minimality of y' we know there is no $w \in \overline{S_i^m}$ such that $r_i(w) \leqslant y'$. Hence we may again just as well write

$$S_i^{m+1} = S_i^m \cup \{w \in \overline{S_i^m} \mid r_i(w) \leqslant y'\}$$

and so we again get $S_i^{m+1} = \mathbf{h}_i^{R,y'}(\vec{S})$. Hence we have shown $\vec{S}^{m+1} = \mathbf{h}^{R,y'}(\vec{S})$. Our result will be proved if we can furthermore show that y' is minimal such that $\vec{S}^m \subset \mathbf{h}^{R,y'}(\vec{S})$. But $y' = \min_{i \in Sources^+}\{r_i(\overline{S_i^m})\}$ implies that, for all $i \in Sources^+$, there is no $w \in \overline{S_i^m}$ such that $r_i(w) < y'$. Hence, for all $i \in Sources^+$ and all $y'' < y'$ we have

$$S_i^m = S_i^m \cup \{w \in \overline{S_i^m} \mid r_i(w) \leqslant y''\} = \mathbf{h}_i^{R,y''}(\vec{S})$$

and so $\vec{S}^m = \mathbf{h}^{R,y''}(\vec{S})$ for all $y'' \leqslant y$. This proves the result. □

Given this lemma, we can now better describe the negotiation process $f^{\mathscr{N}_1(R)}(\vec{S}) = (\vec{S}^0, \ldots, \vec{S}^k)$ on $\vec{S}$ under $\mathscr{N}_1(R)$. The process begins with $\vec{S}^0 = \vec{S} = \mathbf{h}^{R,0}(\vec{S})$. If this is consistent then the process ends, otherwise we carry on . By Lemma 28 we know that $\vec{S}^1 = \mathbf{h}^{R,x_1}(\vec{S})$ where x_1 is minimal such that $\mathbf{h}^{R,0}(\vec{S}) \subset \mathbf{h}^{R,x_1}(\vec{S})$. If this is consistent then we stop, otherwise we carry on. Continuing the process, we see that $f^{\mathscr{N}_1(R)}(\vec{S})$ will take form

$$f^{\mathscr{N}_1(R)}(\vec{S}) = (\mathbf{h}^{R,x_0}(\vec{S}), \mathbf{h}^{R,x_1}(\vec{S}), \ldots, \mathbf{h}^{R,x_k}(\vec{S})),$$

where (i) $x_0 = 0$, (ii) for each $0 \leqslant j$, x_{j+1} is minimal such that $\mathbf{h}^{R,x_j}(\vec{S}) \subset \mathbf{h}^{R,x_{j+1}}(\vec{S})$ and (iii) k is minimal such that $\mathbf{h}^{R,x_k}(\vec{S})$ is consistent. Thus we end up with $\mathbf{f}^{\mathscr{N}_1(R)}(\vec{S}) = \mathbf{h}^{R,x_k}(\vec{S})$. Now, it should be clear that for *all* $x < x_k$ it must be the case that $\mathbf{h}^{R,x}(\vec{S}) = \mathbf{h}^{R,x_j}(\vec{S})$ for some $j < k$. Hence we may state the following corollary to Lemma 28.

Corollary 29. *Let R be a ranking assignment. Then, for all $\vec{S} \in \mathscr{B}^{Sources}$, we have $\mathbf{f}^{\mathscr{N}_1(R)}(\vec{S}) = \mathbf{h}^{R,z}(\vec{S})$, where z is minimal such that $\mathbf{h}^{R,z}(\vec{S})$ is consistent.*

We will now make use of this characterisation of $\mathbf{f}^{\mathscr{N}_1(R)}$ in proving the rest of our results. Before we start we give one more lemma, which as well as being used in proving the next proposition will also be used in the proof of Proposition 25.

Lemma 30. *Let R be a ranking assignment, $\vec{S} \in \mathscr{B}^{Sources}$, $i \in Sources^+$ and $T \in \mathscr{B}$. Let z be minimal such that $\mathbf{h}_i^{R,z}(\vec{S}) \cap T \neq \emptyset$. Then $z = [R_i(S_i)](T)$.*

Proof. Recall that $[R_i(S_i)](T) = \min_{w \in T}[R_i(S_i)](w)$. Clearly we have that there is some $w \in T$ such that $[R_i(S_i)](w) = [R_i(S_i)](T)$, hence we know $\mathbf{h}_i^{R,[R_i(S_i)](T)}(\vec{S}) \cap T \neq \emptyset$, while also we have $[R_i(S_i)](T) \leqslant [R_i(S_i)](w)$ for all $w \in T$. Hence, for all $l < [R_i(S_i)](T)$, there is no $w \in T$ such that $[R_i(S_i)](w) \leqslant l$, i.e., $\mathbf{h}_i^{R,l}(\vec{S}) \cap T = \emptyset$. Hence $[R_i(S_i)](T)$ has the required minimality. □

Proof of Proposition 19. Let $\vec{U} \in \mathscr{B}^{Sources}$ be such that $U_i = S$, $U_0 = T$ and $U_j = \mathscr{W}$ for all $j \notin \{0, i\}$. Then $S \ominus_i^{\mathscr{N}_1(R)} T = \mathbf{f}_i^{\mathscr{N}_1(R)}(\vec{U})$. By Corollary 29 we have $S \ominus_i^{\mathscr{N}_1(R)} T = \mathbf{h}_i^{R,z}(\vec{U})$ where z is minimal such that $\mathbf{h}^{R,z}(\vec{U})$ is consistent. Since $\mathbf{h}_0^{R,z}(\vec{U}) = T$, while clearly $\mathbf{h}_j^{R,z}(\vec{U}) = \mathscr{W}$ for all $j \notin \{0, i\}$, this amounts to saying that z is minimal such that $\mathbf{h}_i^{R,z}(\vec{U}) \cap T \neq \emptyset$. By Lemma 30 we know $z = [R_i(S)](T)$. Hence $S \ominus_i^{\mathscr{N}_1(R)} T = \mathbf{h}_i^{R,[R_i(S)](T)}(\vec{U}) = \{w \in \mathscr{W} \mid [R_i(S)](w) \leqslant [R_i(S)](T)\}$, which completes the proof. □

Proof of Proposition 20. To show (ind4) suppose $w \in S\ominus_i^{\mathcal{N}_1(R)}(T_1 \cup T_2)$. Then, by Proposition 19, $[R_i(S)](w) \leqslant [R_i(S)](T_1 \cup T_2)$. Now, for *any* ranking r and any $A, B \in \mathscr{B}$ such that $A \subseteq B$ it is easy to show that $r(B) \leqslant r(A)$. In particular we have $[R_i(S)](T_1 \cup T_2) \leqslant [R_i(S)](T_1)$ and so we get $[R_i(S)](w) \leqslant [R_i(S)](T_1)$. Using Proposition 19 again this gives us $w \in S\ominus_i^{\mathcal{N}_1(R)} T_1$ and so we have shown $S\ominus_i^{\mathcal{N}_1(R)}(T_1 \cup T_2) \subseteq S\ominus_i^{\mathcal{N}_1(R)} T_1$ as required.

For (ind5) first note that if $(S\ominus_i^{\mathcal{N}_1(R)}(T_1 \cup T_2)) \cap T_1 \neq \emptyset$ then there must exist some $w' \in T_1$ such that $[R_i(S)](w') \leqslant [R_i(S)](T_1 \cup T_2)$. Hence $[R_i(S)](T_1) = \min_{w \in T_1}[R_i(S)](w) \leqslant [R_i(S)](T_1 \cup T_2)$. Thus, making use of Proposition 19, we have that $w \in S\ominus_i^{\mathcal{N}_1(R)} T_1$ implies $[R_i(S)](w) \leqslant [R_i(S)](T_1)$ implies $[R_i(S)](w) \leqslant [R_i(S)](T_1 \cup T_2)$ implies $w \in S\ominus_i^{\mathcal{N}_1(R)}(T_1 \cup T_2)$. Hence we get $S\ominus_i^{\mathcal{N}_1(R)} T_1 \subseteq S\ominus_i^{\mathcal{N}_1(R)}(T_1 \cup T_2)$ as required. □

For the proof of Proposition 21 we will make use of the following lemma.

Lemma 31. *Let R be a ranking assignment and let $x \in \mathbb{N}$. Then, for each $\vec{S} \in \mathscr{B}^{Sources}$, we have* $\bigcap_i \mathbf{h}_i^{R,x}(\vec{S}) = \{w \in S_0 \mid \max_{i \in Sources^+}[R_i(S_i)](w) \leqslant x\}$.

Proof. For this proof let $r(w)$ abbreviate $\max_{i \in Sources^+}[R_i(S_i)](w)$ for each $w \in \mathscr{W}$. Given any $w \in \mathscr{W}$ we have that $w \in \bigcap_i \mathbf{h}_i^{R,x}(\vec{S})$ iff $w \in S_0$ (since $\mathbf{h}_0^{R,x}(\vec{S}) = S_0$) and $[R_i(S_i)](w) \leqslant x$ for all $i \in Sources^+$ (by definition of $\mathbf{h}_i^{R,x}(\vec{S})$ for $i \in Sources^+$). Since saying that $[R_i(S_i)](w) \leqslant x$ for all $i \in Sources^+$ is the same as saying $r(w) \leqslant x$, the result follows. □

Notice that, as a corollary of this result, for each $w \in S_0$ and letting $r(w)$ abbreviate $\max_{i \in Sources^+}[R_i(S_i)](w)$, we always have $w \in \bigcap_i \mathbf{h}_i^{R,r(w)}(\vec{S})$. This fact will be used in the next proof.

Proof of Proposition 21. Let $\vec{S} \in \mathscr{B}^{Sources}$. Then, by definition, $\Delta_{\mathcal{N}_1(R)}(\vec{S}) = \bigcap_i \mathbf{f}_i^{\mathcal{N}_1(R)}(\vec{S})$. By Corollary 29 we know that $\mathbf{f}^{\mathcal{N}_1(R)}(\vec{S}) = \mathbf{h}^{R,z}(\vec{S})$ where z is minimal such that $\mathbf{h}^{R,z}(\vec{S})$ is consistent. Hence $\Delta_{\mathcal{N}_1(R)}(\vec{S}) = \bigcap_i \mathbf{h}_i^{R,z}(\vec{S})$. Applying Lemma 31, then, and again letting $r(w)$ abbreviate $\max_{i \in Sources^+}[R_i(S_i)](w)$ for each $w \in \mathscr{W}$, we get that $\Delta_{\mathcal{N}_1(R)}(\vec{S}) = \{w \in S_0 \mid r(w) \leqslant z\}$. Hence we need to show

$$\{w \in S_0 \mid r(w) \leqslant z\} = \{w \in S_0 \mid r(w) \text{ is minimal}\}.$$

To show this we need to show that, for all $w \in S_0$, $r(w) \leqslant z$ iff $r(w) \leqslant r(w')$ for all $w' \in S_0$. So suppose $w \in S_0$ and that $r(w) \leqslant z$. We claim that, for all $w' \in S_0$, we have $z \leqslant r(w')$. To see this, suppose $w' \in S_0$ was such that $r(w') < z$. By the remark following Lemma 31 we have that $w' \in \bigcap_i \mathbf{h}_i^{R,r(w')}(\vec{S})$ and so $\bigcap_i \mathbf{h}_i^{R,r(w')}(\vec{S}) \neq \emptyset$. But this contradicts the minimality of z, and so we must have $z \leqslant r(w')$ for all $w' \in S_0$ as claimed. Given this we can deduce from $r(w) \leqslant z$ that $r(w) \leqslant r(w')$ for all $w' \in S_0$ as required. For the converse direction, let $w \in S_0$ be such that $r(w) \leqslant r(w')$ for all $w' \in S_0$. Since $\mathbf{h}^{R,z}(\vec{S})$ is consistent, we know that there exists some $w_0 \in \bigcap_i \mathbf{h}_i^{R,z}(\vec{S})$, i.e., that there exists some $w_0 \in S_0$ such that $r(w_0) \leqslant z$. Hence in particular we get $r(w) \leqslant r(w_0) \leqslant z$ as required. This completes the proof. □

Proof of Proposition 23. Let R be a ranking assignment, $\vec{S} \in \mathscr{B}^{Sources}$ and $i \in Sources^+$. For this proof, we will use r_i to denote the ranking $[R_i(S_i)]$. By Corollary 29 we know that $\mathbf{f}^{\mathcal{N}_1(R)}(\vec{S}) = \mathbf{h}^{R,z}(\vec{S})$ where z is minimal such that $\mathbf{h}^{R,z}(\vec{S})$ is consistent. We first claim that $r_i\left(\bigcap_{j \neq i} \mathbf{f}_j^{\mathcal{N}_1(R)}(\vec{S})\right) \leqslant z$. To see this, note that, since $\mathbf{h}^{R,z}(\vec{S})$ is consistent, we know there exists some $w' \in \mathscr{W}$ such that $w' \in \bigcap_j \mathbf{h}_j^{R,z}(\vec{S})$. In particular we have $w' \in \mathbf{h}_i^{R,z}(\vec{S})$ and so $r_i(w') \leqslant z$. Meanwhile, since also $w' \in \bigcap_{j \neq i} \mathbf{h}_j^{R,z}(\vec{S}) = \bigcap_{j \neq i} \mathbf{f}_j^{\mathcal{N}_1(R)}(\vec{S})$, we have $r_i\left(\bigcap_{j \neq i} \mathbf{f}_j^{\mathcal{N}_1(R)}(\vec{S})\right) \leqslant r_i(w')$. Putting these two inequalities together gives us $r_i\left(\bigcap_{j \neq i} \mathbf{f}_j^{\mathcal{N}_1(R)}(\vec{S})\right) \leqslant z$ as claimed. Now to prove the proposition, let $w \in \mathscr{W}$ be such that $w \in S_i\ominus_i^{\mathcal{N}_1(R)}\left(\bigcap_{j \neq i} \mathbf{f}_j^{\mathcal{N}_1(R)}(\vec{S})\right)$. Then, by Proposition 19, $r_i(w) \leqslant r_i\left(\bigcap_{j \neq i} \mathbf{f}_j^{\mathcal{N}_1(R)}(\vec{S})\right)$. Hence $r_i(w) \leqslant z$, equivalently $w \in \mathbf{h}_i^{R,z}(\vec{S}) = \mathbf{f}_i^{\mathcal{N}_1(R)}(\vec{S})$. Thus we have shown $S_i\ominus_i^{\mathcal{N}_1(R)}\left(\bigcap_{j \neq i} \mathbf{f}_j^{\mathcal{N}_1(R)}(\vec{S})\right) \subseteq \mathbf{f}_i^{\mathcal{N}_1(R)}(\vec{S})$ as required. □

A.4. Proofs from Section 4.2

Proof of Proposition 25. Suppose $f^{\mathcal{N}_2(R)}(\vec{S}) = \sigma = (\vec{S}^0, \ldots, \vec{S}^k)$ and let $i \in Sources^+$. In this proof we will denote i's ranking $[R_i(S_i)]$ by just r_i. Let l be minimal such that $\bigcap_{j<i} S_j^l \neq \emptyset$. In other words, σ_l is that stage in the negotiation where it is i's turn to weaken. Note that i has not done any weakening up to this stage, i.e., $S_i^l = S_i^0 = S_i$. This is because if i *had* already weakened, i.e., we had $g_2(\sigma_t) = \{i\}$ for some $t < l$, then by definition of g_2 this would mean $\bigcap_{j<i} S_j^t \neq \emptyset$—contradicting the minimality of l. Also, let l' be minimal such that $\bigcap_{j \leqslant i} S_j^{l'} \neq \emptyset$. Then between stages σ_l and $\sigma_{l'}$, source i—and only source i—is required to weaken, i.e., we have $g_2(\sigma_s) = \{i\}$ for all $l \leqslant s < l'$. We now make use again of the $\mathbf{h}$-notation which we introduced in the last section. First we need the following lemma.

Lemma 32. *For each $l \leqslant s < l'$, if $S_i^s = \mathbf{h}_i^{R,x}(\vec{S})$ for some x, then $S_i^{s+1} = \mathbf{h}_i^{R,y}(\vec{S})$, where y is minimal such that $\mathbf{h}_i^{R,x}(\vec{S}) \subset \mathbf{h}_i^{R,y}(\vec{S})$.*

Proof. Let $l \leqslant s < l'$ and suppose $S_i^s = \mathbf{h}_i^{R,x}(\vec{S})$ for some x. Since $g_2(\sigma_s) = \{i\}$ we have $S_i^{s+1} = \blacktriangledown_{\sigma_s}(S_i^s) = S_i^s \cup \{w \in \overline{S_i^s} \mid r_i(w) \text{ is minimal}\}$. Let $y = \min_{w \in \overline{S_i^s}} r_i(w)$. We claim then that $S_i^{s+1} = \mathbf{h}_i^{R,y}(\vec{S})$. To see this note that, using the fact that $S_i^s = \mathbf{h}_i^{R,x}(\vec{S})$, we may re-express S_i^{s+1} as

$$\begin{aligned} S_i^{s+1} &= \{w \mid r_i(w) \leqslant x\} \cup \{w \mid r_i(w) = y\} \\ &= \{w \mid r_i(w) \leqslant y\} \text{ using the minimality of } y \\ &= \mathbf{h}_i^{R,y}(\vec{S}) \text{ as claimed.} \end{aligned}$$

It remains to show that y is in fact minimal such that $\mathbf{h}_i^{R,x}(\vec{S}) \subset \mathbf{h}_i^{R,y}(\vec{S})$. But, by the minimality of y, for each $x \leqslant y' < y$ we know there is no $w \in \overline{S_i^s}$ such that $r_i(w) \leqslant y'$. Hence, for all $x \leqslant y' < y$ we have $\mathbf{h}_i^{R,x}(\vec{S}) = \mathbf{h}_i^{R,y'}(\vec{S})$ and so y is indeed minimal such that $\mathbf{h}_i^{R,x}(\vec{S}) \subset \mathbf{h}_i^{R,y}(\vec{S})$. □

Using this lemma, we can now see that, for each $l \leqslant s \leqslant l'$, we have $S_i^s = \mathbf{h}_i^{R,x_s}(\vec{S})$ where (i) $x_l = 0$ (since $S_i^l = S_i$) and (ii) for $l \leqslant s < l'$, x_{s+1} is minimal such that $\mathbf{h}_i^{R,x_s}(\vec{S}) \subset \mathbf{h}_i^{R,x_{s+1}}(\vec{S})$. In particular we have $S_i^{l'} = \mathbf{h}_i^{R,x_{l'}}(\vec{S})$. Now, since we assumed $\bigcap_{j \leqslant i} S_j^{l'} \neq \emptyset$, we know $\bigcap_{j<i} S_j^{l'} \cap \mathbf{h}_i^{R,x_{l'}}(\vec{S}) \neq \emptyset$. We now claim that, for *all* $t < x_{l'}$ we have $\bigcap_{j<i} S_j^{l'} \cap \mathbf{h}_i^{R,t}(\vec{S}) = \emptyset$. This follows since if $t < x_{l'}$ then we must have $\mathbf{h}_i^{R,t}(\vec{S}) = \mathbf{h}_i^{R,x_s}(\vec{S})$ for some $l \leqslant s < l'$, i.e., $\mathbf{h}_i^{R,t}(\vec{S}) = S_i^s$ for some $l \leqslant s < l'$, and by the minimality of l' we know $\bigcap_{j \leqslant i} S_j^s = \emptyset$. Since $S_j^{l'} = S_j^s$ for all $j \in Sources$ such that $j \neq i$ this gives $\bigcap_{j<i} S_j^{l'} \cap \mathbf{h}_i^{R,t}(\vec{S}) = \emptyset$ as claimed. Hence we have shown in fact that

$$S_i^{l'} = \mathbf{h}_i^{R,z}(\vec{S})$$

where z is minimal such that $\bigcap_{j<i} S_j^{l'} \cap \mathbf{h}_i^{R,z}(\vec{S}) \neq \emptyset$. Now, by Lemma 30, we know that $z = r_i\left(\bigcap_{j<i} S_j^{l'}\right)$. Hence we have

$$S_i^{l'} = \left\{ w \in \mathscr{W} \mid r_i(w) \leqslant r_i\left(\bigcap_{j<i} S_j^{l'}\right)\right\}.$$

Now, by definition of g_2 we know that, after stage $\sigma_{l'}$, neither source i nor any of the sources $j < i$ do any further weakening. This is because $\bigcap_{j \leqslant i} S_j^{l'} \neq \emptyset$ and so, for all $l' < s$ and all sources $j' \leqslant i$ we will have $\bigcap_{j \leqslant j'} S_j^s \neq \emptyset$ (since $\vec{S}^{l'} \subseteq \vec{S}^s$). Hence we know $\mathbf{f}_j^{\mathscr{N}_2(R)}(\vec{S}) = S_j^k = S_j^{l'}$ for all $j \leqslant i$. Hence we get

$$\begin{aligned} \mathbf{f}_i^{\mathscr{N}_2(R)}(\vec{S}) = S_i^{l'} &= \left\{ w \in \mathscr{W} \mid r_i(w) \leqslant r_i\left(\bigcap_{j<i} S_j^{l'}\right)\right\} \\ &= \left\{ w \in \mathscr{W} \mid r_i(w) \leqslant r_i\left(\bigcap_{j<i} \mathbf{f}_j^{\mathscr{N}_2(R)}(\vec{S})\right)\right\} \\ &= S_i \ominus_i^{\mathscr{N}_2(R)} \left(\bigcap_{j<i} \mathbf{f}_j^{\mathscr{N}_2(R)}(\vec{S})\right). \end{aligned}$$

This last step follows from Propositions 19 and 14. This completes the proof of Proposition 25. □

Proof of Proposition 26. We need to show that, for all $i \in Sources^+$,

$$\mathbf{f}_i^{\mathscr{N}_2(R)}(\vec{S}) = S_i \ominus_i^{\mathscr{N}_2(R)} \left(\bigcap_{j \neq i} \mathbf{f}_j^{\mathscr{N}_2(R)}(\vec{S})\right).$$

Letting $X = \bigcap_{j<i} \mathbf{f}_j^{\mathscr{N}_2(R)}(\vec{S})$ and $Y = \bigcap_{i<j} \mathbf{f}_j^{\mathscr{N}_2(R)}(\vec{S})$ this means we must show $\mathbf{f}_i^{\mathscr{N}_2(R)}(\vec{S}) = S_i \ominus_i^{\mathscr{N}_2(R)}(X \cap Y)$. Proposition 25 tells us $\mathbf{f}_i^{\mathscr{N}_2(R)}(\vec{S}) = S_i \ominus_i^{\mathscr{N}_2(R)} X = S_i \ominus_i^{\mathscr{N}_2(R)} (X \cup (X \cap Y))$, hence it suffices to show $S_i \ominus_i^{\mathscr{N}_2(R)}(X \cup (X \cap Y)) = S_i \ominus_i^{\mathscr{N}_2(R)}(X \cap Y)$. Using the fact that $\ominus_i^{\mathscr{N}_2(R)}$ satisfies (ind4) and (ind5) from Proposition 20 we know that this equality holds, provided that $(S_i \ominus_i^{\mathscr{N}_2(R)} X) \cap (X \cap Y) \neq \emptyset$. But $S_i \ominus_i^{\mathscr{N}_2(R)} X = \mathbf{f}_i^{\mathscr{N}_2(R)}(\vec{S})$ by Proposition 25, while $X \cap Y = \bigcap_{j \neq i} \mathbf{f}_j^{\mathscr{N}_2(R)}(\vec{S})$. Hence

$$(S_i \ominus_i^{\mathscr{N}_2(R)} X) \cap (X \cap Y) = \bigcap_{j \in Sources} \mathbf{f}_j^{\mathscr{N}_2(R)}(\vec{S})$$

and this is non-empty by (sc2), as required. □

Proof of Proposition 27. To improve readability, let us denote $\mathbf{f}^{\mathscr{N}_2(R)}$ by just $\mathbf{f}$ in this proof. Given tuples $\vec{x}, \vec{y} \in \mathbb{N}^n$, we will write $\vec{x} \leqslant_{lex} \vec{y}$ whenever either $\vec{x} <_{lex} \vec{y}$ or $\vec{x} = \vec{y}$.

"$\subseteq$". Let $w \in \Delta_{\mathscr{N}_2(R)}(\vec{S})$, i.e., $w \in \bigcap_i \mathbf{f}_i(\vec{S})$. We must show (i) $w \in S_0$, and (ii) for all $w' \in S_0$, we have $(r_1(w), \ldots, r_n(w)) \leqslant_{lex} (r_1(w'), \ldots, r_n(w'))$. Since $w \in \mathbf{f}_0(\vec{S}) = S_0$ we know (i) holds. To show (ii) let $w' \in S_0$. If $(r_1(w'), \ldots, r_n(w')) = (r_1(w), \ldots, r_n(w))$ then we are done, so suppose instead $(r_1(w'), \ldots, r_n(w')) \neq (r_1(w), \ldots, r_n(w))$ and let j be minimal such that $r_j(w') \neq r_j(w)$. We must show $r_j(w) < r_j(w')$.But since $w \in \bigcap_i \mathbf{f}_i(\vec{S})$ we have $w \in \mathbf{f}_j(\vec{S})$ and so, by Propositions 25 and 19, we have $r_j(w) \leqslant r_j\left(\bigcap_{k<j} \mathbf{f}_k(\vec{S})\right)$, equivalently $r_j(w) \leqslant r_j(w'')$ for all $w'' \in \bigcap_{k<j} \mathbf{f}_k(\vec{S})$. Hence if we could show $w' \in \bigcap_{k<j} \mathbf{f}_k(\vec{S})$ then we would get $r_j(w) \leqslant r_j(w')$ and so, since $r_j(w') \neq r_j(w)$, this would give the required $r_j(w) < r_j(w')$. So let $k < j$. For $k = 0$ we already know $w' \in S_0 = \mathbf{f}_0(\vec{S})$. So assume $k \neq 0$. Then by Propositions 25 and 19, to show $w' \in \mathbf{f}_k(\vec{S})$ we need to show $r_k(w') \leqslant r_k\left(\bigcap_{s<k} \mathbf{f}_s(\vec{S})\right)$. But, by the minimality of j we have $r_k(w') = r_k(w)$ and, since $w \in \mathbf{f}_k(\vec{S})$, we have $r_k(w) \leqslant r_k\left(\bigcap_{s<k} \mathbf{f}_s(\vec{S})\right)$. This gives the required conclusion. Hence $w' \in \bigcap_{k<j} \mathbf{f}_k(\vec{S})$ as required.

"$\subseteq$". Let $w \in S_0$ be such that $(r_1(w), \ldots, r_n(w)) \leqslant_{lex} (r_1(w'), \ldots, r_n(w'))$ for all $w' \in S_0$. We must show $w \in \bigcap_i \mathbf{f}_i(\vec{S})$. We already have $w \in \mathbf{f}_0(\vec{S}) = S_0$. We will now show by induction on k that $w \in \mathbf{f}_k(\vec{S})$ for each $k = 1, \ldots, n$. For $k = 1$ we need to show (by Propositions 25 and 19) that $r_1(w) \leqslant r_1(w')$ for all $w' \in S_0$. But this follows already from the fact that $(r_1(w), \ldots, r_n(w)) \leqslant_{lex} (r_1(w'), \ldots, r_n(w'))$ for all $w' \in S_0$.Now let $1 < k \leqslant n$ and suppose for induction that $w \in \bigcap_{s<k} \mathbf{f}_s(\vec{S})$. By Propositions 25 and 19 we need to

show $r_k(w) \leqslant r_k(w')$ for all $w' \in \bigcap_{s<k}\mathbf{f}_s(\vec{S})$. So let $w' \in \bigcap_{s<k}\mathbf{f}_s(\vec{S})$. Since $w' \in \mathbf{f}_0(\vec{S}) = S_0$, we know $(r_1(w),\ldots,r_n(w)) \leqslant_{lex} (r_1(w'),\ldots,r_n(w'))$. Hence our result will be proved if $r_s(w) = r_s(w')$ for all $s < k$. But, given $s < k$, we have $w' \in \mathbf{f}_s(\vec{S})$ and so $r_s(w') \leqslant r_s(w'')$ for all $w'' \in \bigcap_{t<s}\mathbf{f}_t(\vec{S})$. Since, using induction, $w \in \bigcap_{t<s}\mathbf{f}_t(\vec{S})$, this means in particular that $r_s(w') \leqslant r_s(w)$. By a symmetric argument it can be shown that $r_s(w) \leqslant r_s(w')$, hence $r_s(w) = r_s(w')$ as required. Hence we have shown $w \in \bigcap_k \mathbf{f}_k(\vec{S})$. This completes the inductive step and so $w \in \bigcap_i \mathbf{f}_i(\vec{S})$ as required. □

References

[1] C. Alchourrón, P. Gärdenfors, D. Makinson, On the logic of theory change: Partial meet contraction and revision functions, Journal of Symbolic Logic 50 (1985) 510–530.

[2] M. Arenas, L. Bertossi, J. Chomicki, Consistent query answers in inconsistent databases, in: Proceedings of the Eighteenth ACM SIGACT SIGMOD SIGART Symposium on Principles of Database Systems (PODS'99), 1999, pp. 68–79.

[3] O. Arieli, B. Van Nuffelen, M. Denecker, M. Bruynooghe, Coherent composition of distributed knowledge-bases through abduction, in: Proceedings of the Eighth International Conference on Logic for Programming, Artificial Intelligence and Reasoning (LPAR'01), LNCS, 2250, Springer, Berlin, 2001, pp. 624–638.

[4] C. Baral, S. Kraus, J. Minker, V.S. Subrahmanian, Combining multiple knowledge bases consisting of first order theories, Computational Intelligence 8 (1992) 45–71.

[5] R. Booth, A negotiation-style framework for non-prioritised revision, in: Proceedings of the Eighth Conference on Theoretical Aspects of Rationality and Knowledge (TARK 2001), 2001, pp. 137–150.

[6] R. Booth, S. Chopra, A. Ghose, T. Meyer, Belief liberation (and retraction), in: Proceedings of the Ninth Conference on Theoretical Aspects of Rationality and Knowledge (TARK 2003), 2003, pp. 159–172.

[7] J. Cantwell, Resolving conflicting information, Journal of Logic, Language and Information 7 (2) (1998) 191–220.

[8] P. Gärdenfors, Knowledge in Flux, MIT Press, Cambridge, MA, 1988.

[9] J. Grant, J. Minker, A logic-based approach to data integration, Theory and Practice of Logic Programming 2 (3) (2002) 323–368.

[10] G. Greco, S. Greco, E. Zumpano, A logic programming approach to the integration, repairing and querying of inconsistent databases, in: Proceedings of the Seventeenth International Conference on Logic Programming (ICLP'01), LNCS, 2237, Springer, Berlin, 2001, pp. 348–364.

[11] S.O. Hansson, A Textbook of Belief Dynamics, Kluwer Academic Publishers, Dordrecht, 1999.

[12] S. Konieczny, R. Pino-Pérez, On the logic of merging, in: Proceedings of the Sixth International Conference on Principles of Knowledge Representation and Reasoning (KR'98), 1998, pp. 488–498.

[13] S. Konieczny, R. Pino-Pérez, Merging information under constraints: a qualitative framework, Journal of Logic and Computation 12 (5) (2002) 773–808.

[14] I. Levi, Subjunctives, dispositions and chances, Synthese 34 (1977) 423–455.

[15] P. Liberatore, M. Schaerf, Arbitration (or how to merge knowledge bases), IEEE Transactions on Knowledge and Engineering 10 (1) (1998) 76–90.

[16] P. Maynard-Reid II, Y. Shoham, Belief fusion: aggregating pedigreed belief states, Journal of Logic, Language and Information 10 (2) (2001) 183–209.

[17] T. Meyer, On the semantics of combination operations, Journal of Applied Non-Classical Logics 11 (1–2) (2001) 59–84.

[18] T. Meyer, N. Foo, R. Kwok, D. Zhang, Logical foundations of negotiation: strategies and preferences, in: Proceedings of the Ninth International Conference on Principles of Knowledge Representation and Reasoning (KR'04), 2004, pp. 311–318.

[19] T. Meyer, A. Ghose, S. Chopra, Social choice, merging and elections, in: Proceedings of the Sixth European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty (ECSQARU 2001), 2001, pp. 466–477.

[20] P. Revesz, On the semantics of theory change: Arbitration between old and new information, in: Proceedings of the Twelfth ACM SIGACT SIGMOD SIGART Symposium on Principles of Database Systems (PODS'93), 1993, pp. 71–82.

[21] H. Rott, M. Pagnucco, Severe withdrawal (and recovery), Journal of Philosophical Logic 28 (1999) 501–547.

[22] W. Spohn, Ordinal conditional functions: A dynamic theory of epistemic states, in: W.L. Harper, B. Skyrms (Eds.), Causation in Decision, Belief Change and Statistics, vol. 2, 1988.

[23] M. A. Williams, Transmutations of knowledge systems, in: Proceedings of the Fourth International Conference on Principles of Knowledge Representation and Reasoning (KR'94), 1994, pp. 619–629.

[24] D. Walton, E. Krabbe, Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning, State University of New York Press, Albany, 1995.

[25] M. Wooldridge, S. Parsons, Languages for negotiation, in: Proceedings of the Fourteenth European Conference on Artificial Intelligence (ECAI 2000), 2000, pp. 393–400.