Social contraction and belief negotiation

Richard Booth
University of Leipzig
Department of Computer Science
Augustusplatz 10/11
04109 Leipzig, Germany
booth@informatik.uni-leipzig.de

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 framework of *belief negotiation models* (Booth 2001). 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.

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, for e.g., (Booth 2001, Cantwell 1998, Konieczny & Pino-Pérez 1998, Konieczny & Pino-Pérez 1999, Liberatore & Schaerf 1998, Maynard-Reid & Shoham 2001, Meyer 2001, Revesz 1993)). 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* (Alchourrón et al. 1985, Gärdenfors 1988, Hansson 1999). 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 (Levi 1977) and is given succinct expression by the *Levi Identity* (Gärdenfors 1988), 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, in which several items of information may all be weakened simultaneously so that they become 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 nego-*

tiation models, which was introduced in (Booth 2001) 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 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.

Preliminaries: We let $\mathcal{W}$ 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 $\mathcal{W}$ we denote by $\mathcal{B}$. Given $S \subseteq \mathcal{W}$, we use $\overline{S}$ to denote $\mathcal{W} - S$. We assume throughout that we have a fixed finite set $Sources = \{1, \ldots, n\}$ of information sources ($n \geq 2$). We work semantically throughout, so each item of information provided by a source i will take the form of a set $S_i \in \mathcal{B}$ (so no source ever provides the "inconsistent" information $\emptyset$). An *information profile* (relative to *Sources*) is an element of $\mathcal{B}^n$. We use $\vec{S}, \vec{S}^1$, etc. to denote information profiles, with S_i, S_i^1, etc. denoting the i^{th} element of $\vec{S}, \vec{S}^1$, etc. The idea is that S_i is the information in $\vec{S}$ belonging to source i. We will say an information profile $\vec{S}$ is *consistent* when $\bigcap_i S_i \neq \emptyset$, 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 $S_i^1 \subseteq S_i^2$ for all $i \in Sources$. Finally if $\mathbf{f}$ is a function with range $\mathcal{B}^n$, we will use $\mathbf{f}_i(\vec{S})$ to denote the i^{th} element of $\mathbf{f}(\vec{S})$.

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 $\mathbf{f} : \mathcal{B}^n \to \mathcal{B}^n$ 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 properties of such an $\mathbf{f}$:

(sc1) $\quad \vec{S} \subseteq \mathbf{f}(\vec{S})$

(sc2) $\quad \mathbf{f}(\vec{S})$ is consistent

(sc3) $\quad$ If $\vec{S}$ is consistent then $\mathbf{f}(\vec{S}) = \vec{S}$

(sc1) decrees that the modification is carried out by *weakening* the individual items of information. **(sc2)** says that the end results of all these weakenings should be jointly consistent. **(sc3)** says that if $\vec{S}$ is already consistent then 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 n to be this completely reliable source**, and reflect this by insisting on the following rule for SC functions:

(sc4) $\quad \mathbf{f}_n(\vec{S}) = S_n$

We will denote the set of sources *minus* n by *Sources**. We make the following definition:

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

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 when that item is considered completely reliable.

Definition 2 *Let* $\mathbf{f}$ *be an SC function and let* $i \in Sources^*$. *We define the function* $\ominus_i^{\mathbf{f}} : \mathcal{B} \times \mathcal{B} \to \mathcal{B}$ *by, for all* $S, T \in \mathcal{B}$, $S \ominus_i^{\mathbf{f}} T = \mathbf{f}_i(\vec{U})$, *where* $\vec{U} \in \mathcal{B}^n$ *is such that* $U_i = S, U_n = T$ *and* $U_j = \mathcal{W}$ *for all* $j \notin \{i, n\}$. *We call* $\ominus_i^{\mathbf{f}}$ *i's* individual contraction function (relative to $\mathbf{f}$).

(E.g., if $n = 4$, then $S \ominus_2^{\mathbf{f}} T$ is the 2nd entry of the 4-tuple $\mathbf{f}(\mathcal{W}, S, \mathcal{W}, T)$.) 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. The following proposition is easy to prove.

Proposition 1 *Let* $\mathbf{f}$ *be an SC function and let* $i \in Sources^*$. *Then* $\ominus_i^{\mathbf{f}}$ *satisfies* **(ic1)** $S \subseteq S \ominus_i^{\mathbf{f}} T$, **(ic2)** $(S \ominus_i^{\mathbf{f}} T) \cap T \neq \emptyset$, *and* **(ic3)** *If* $S \cap T \neq \emptyset$ *then* $S \ominus_i^{\mathbf{f}} T = S$.

The properties **(ic1)**–**(ic3)** essentially correspond to the well-known basic AGM postulates for contraction (Alchourrón et al. 1985) *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 $\ominus_i^{\mathbf{f}}$ don't generally satisfy this much debated (see (Hansson 1999 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. 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 $\vec{S}$,

$$\Delta_{\mathbf{f}}(\vec{S}) = \bigcap_{i=1}^{n} \mathbf{f}_i(\vec{S}).$$

Our basic postulates for $\mathbf{f}$ 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 $\Delta_{\mathbf{f}}(\vec{S}) \subseteq S_n$, i.e., the result of the merging must always imply the information provided by source n. 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 operators for short (Konieczny & Pino-Pérez 1999), S_n here taking the role of the integrity constraints.[1]

[1] 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 (Gärdenfors 1988). 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. 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 $\mathbf{f}_n(\vec{S}) = S_n$ 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)**. It is also quite easy to construct simple counter-examples which show that **(sc5)** doesn't hold in general for SC functions. (E.g. define $\mathbf{f}$ by setting, for each $i \in Sources$, $\mathbf{f}_i(\vec{S}) = S_i$ if either $i = n$ or $\vec{S}$ is consistent, $\mathbf{f}_i(\vec{S}) = \mathcal{W}$ otherwise.) 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 1 Suppose we have three sources, i.e., $n = 3$. Suppose source 1 provides the information $S \neq \mathcal{W}$, source 2 provides the complete opposite information $\overline{S}$, and the completely reliable source 3 provides only the trivial information $\mathcal{W}$. We first claim that for any SC function $\mathbf{f}$ relative to these sources which satisfies **(sc5)** we have either $\mathbf{f}_1(S, \overline{S}, \mathcal{W}) = S$ or $\mathbf{f}_2(S, \overline{S}, \mathcal{W}) = \overline{S}$. To see this, suppose $\mathbf{f}_1(S, \overline{S}, \mathcal{W}) \neq S$. Then, by **(sc5)**, we must have $S \cap \mathbf{f}_2(S, \overline{S}, \mathcal{W}) \cap \mathbf{f}_3(S, \overline{S}, \mathcal{W}) = \emptyset$. Now we know by **(sc4)** (or **(sc1)**) that $\mathbf{f}_3(S, \overline{S}, \mathcal{W}) = \mathcal{W}$. Hence we have $S \cap \mathbf{f}_2(S, \overline{S}, \mathcal{W}) = \emptyset$, i.e., $\mathbf{f}_2(S, \overline{S}, \mathcal{W}) \subseteq \overline{S}$. Since we also have $\overline{S} \subseteq \mathbf{f}_2(S, \overline{S}, \mathcal{W})$ by **(sc1)**, we conclude that $\mathbf{f}_2(S, \overline{S}, \mathcal{W}) = \overline{S}$ which proves the claim. Given this, we have for the corresponding merging operator that either $\Delta_{\mathbf{f}}(S, \overline{S}, \mathcal{W}) \subseteq S$ or $\Delta_{\mathbf{f}}(S, \overline{S}, \mathcal{W}) \subseteq \overline{S}$. 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 with-hold judgement on whether S or $\overline{S}$ holds in the merging and to expect, say, $\Delta_{\mathbf{f}}(S, \overline{S}, \mathcal{W}) = \mathcal{W}$. 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 (Hansson 1999 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 n) 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" (Hansson 1999 p. 77) which, in turn, is a characteristic postulate of maxichoice contraction:

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

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 $\mathbf{f}_i(\vec{S}) = \mathcal{W}$,

The rule **(sc1+)** says, in effect, that the information from each source is either kept or discarded completely.[2]

Although Example 1 suggests **(sc5)** may be too strong, possible weakenings of it are at hand. One, which brings the individual contraction functions into the picture, 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 $\mathbf{f}_i(\vec{S}) \subseteq S_i \ominus_i^{\mathbf{f}} \overline{S_i}$

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 **(sc5–)** 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.

The last postulate we look at is motivated by the feeling that social contraction should be entirely expressible in terms of the individual contraction functions.

(sc6) For all $i \in Sources^*$,
$\mathbf{f}_i(\vec{S}) = S_i \ominus_i^{\mathbf{f}} (\bigcap_{j \neq i} \mathbf{f}_j(\vec{S}))$

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 1, $\ominus_i^{\mathbf{f}}$ satisfies **(ic3)**, it is easy to see that any SC function satisfying **(sc6)** also satisfies **(sc5)**. (In fact only the "$\subseteq$" direction of **(sc6)** is needed to prove **(sc5)**.)

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 (Booth 2001) 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 different sources (one of which is considered completely reliable) and show how each such extended belief negotiation model $\mathcal{N}$ yields an SC function $\mathbf{f}^{\mathcal{N}}$. Let's begin with a rough description of the framework.[3]

Suppose the information profile $\vec{S}$ is provided by *Sources*. The idea is that we determine $\mathbf{f}^{\mathcal{N}}(\vec{S})$ as follows. We start off with the information profile $\vec{S}^0 = \vec{S}$. If $\vec{S}^0$ is consistent then we just take $\mathbf{f}^{\mathcal{N}}(\vec{S}) = \vec{S}^0$. 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 con-

[2] Precisely such an assumption is made explicitly in (Cantwell 1998). Its adoption here would effectively reduce social contraction to something akin to *belief base contraction* (Hansson 1999).

[3] We remark that this framework shares some similarities with the abstract formalisation of negotiation found in (Wooldridge & Parsons 2000). For a more detailed treatment of negotiation see (Walton & Krabbe 1995).

cessions", 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 $\mathbf{f}^{\mathcal{N}}(\vec{S}) = \vec{S}^1$. 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 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 n is completely reliable (and so never loses a round) by taking the range of g to be $2^{Sources^*}$ rather than $2^{Sources}$. We make two more mild restrictions on g. Firstly, in order to avoid deadlock we need to assume that at least one party must weaken at each stage:

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

Secondly, 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) $\quad i \in g(\sigma)$ 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 given that i must 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) $\quad S_i^m \subseteq \blacktriangledown_\sigma(i)$

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

We can now make the following definition.

Definition 3 *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.[4]

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 $f^{\mathcal{N}} : \mathcal{B}^n \to \Omega$ 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 $0 \leq j < k$ we have, for each $i \in Sources$,

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

A belief negotiation model $\mathcal{N}$ thus yields a function $\mathbf{f}^{\mathcal{N}} : \mathcal{B}^n \to \mathcal{B}^n$, via $f^{\mathcal{N}}$ above, by simply taking $\mathbf{f}^{\mathcal{N}}(\vec{S}) = \vec{S}^k$. It is easy to check that $\mathbf{f}^{\mathcal{N}}$ forms an SC function, and not much harder to show that, in fact, *every* SC function can be said to arise in this way:

Theorem 1 *Let $\mathbf{f} : \mathcal{B}^n \to \mathcal{B}^n$ be a function. Then $\mathbf{f}$ is an SC function iff $\mathbf{f} = \mathbf{f}^{\mathcal{N}}$ for some belief negotiation model $\mathcal{N}$.*

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

[4] There are a couple of slight notational differences between this paper and (Booth 2001). 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.

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 2 *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 $S_i \in \mathcal{B}$, but also some indication of the *plausibility* of *all* the worlds in $\mathcal{W}$. Such an indication is provided by a *ranking.*

Definition 4 *A* ranking *is a function* $r : \mathcal{W} \to \mathbb{N}$. *We extend such an* r *to a function on* $2^{\mathcal{B}}$ *by setting, for each* $T \in \mathcal{B}$, $r(T) = \min_{w\in T} r(w)$. *Given* $S \in \mathcal{B}$ *we say that* r *is a ranking* relative to S *iff* $r^{-1}(0) = S$.

Such rankings, or variants thereof, are a popular tool in knowledge representation. They can be traced back to the work of (Spohn 1988) and indeed have already been employed in the context of both merging (see e.g. (Meyer 2001, Meyer et al. 2001)) and belief revision (see e.g. (Williams 1994)). A ranking provides, for each $w \in \mathcal{W}$, a measure of the plausibility of w being the actual world. The lower $r(w)$ is, the more plausible it is considered to be. Rankings also allow us to talk about *degrees of certainty* or *belief.* Given $S \in \mathcal{B}$, we can interpret $r(\overline{S})$ as the degree of certainty that the world is in S — the higher $r(\overline{S})$ 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 relative to S_i. Formally, we assume we are given a *ranking assignment* for *Sources*:

Definition 5 *A* ranking assignment *(relative to Sources) is a function* R *which assigns, to each* $i \in Sources^*$ *and* $S \in \mathcal{B}$, *a ranking* $[R_i(S)]$ *relative to* S.

Note we assume source n does not provide a ranking, just S_n as normal.[5] 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 which they have provided along with their initial information S_i^0. More precisely we define g_1 from R by setting

$$g_1(\sigma) \quad = \quad \{i \in Sources^* \mid S_i^m \neq \mathcal{W} \text{ and } [R_i(S_i^0)](\overline{S_i^m}) \text{ is minimal}\}$$

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

$$\blacktriangledown_\sigma(i) = S_i^m \cup \{w \in \overline{S_i^m} \mid [R_i(S_i^0)](w) \text{ is minimal}\}$$

Given a ranking assignment R, we let $\mathcal{N}(R)$ 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 3.) Let's now see an example of $\mathcal{N}(R)$ "in action".

Example 2 For this example we assume our background propositional language contains just two propositional variables, leading $\mathcal{W}$ to contain just four worlds which we denote here by a, b, c, d. We also assume that $Sources = \{1, 2, 3\}$. Suppose source 1 gives initial information $\{a\}$, source 2 gives $\{c\}$ and completely reliable source 3 gives $\mathcal{W}$ (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 given as follows:

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

Here, the columns correspond to ranks. So, for example, $[R_1(\{a\})]$ gives world b a rank of 1, c a

[5] We also make an assumption of *commensurability* (Meyer et al. 2001), i.e., that all sources use the same *scale* when ranking the worlds according to plausibility.

rank of 2 etc. We construct the complete negotiation process $f^{\mathcal{N}(R)}(\{a\},\{c\},\mathcal{W}) = \sigma$ stage by stage, starting with $\sigma_0 = (\langle\{a\},\{c\},\mathcal{W}\rangle)$. 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 $\blacktriangledown_{\sigma_0}(1) = \{a\} \cup \{w \in \overline{\{a\}} \mid [R_1(\{a\})](w) \text{ is minimal}\}$, 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 $\sigma_1 = (\langle\{a\},\{c\},\mathcal{W}\rangle, \langle\{a,b\},\{c\},\mathcal{W}\rangle)$. 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 $\blacktriangledown_{\sigma_1}(1) = \{a,b\} \cup \{w \in \overline{\{a,b\}} \mid [R_1(\{a\})](w) \text{ is minimal}\} = \{a,b,c,d\} = \mathcal{W}$ and $\blacktriangledown_{\sigma_1}(2) = \{c\} \cup \{w \in \overline{\{c\}} \mid [R_2(\{c\})](w) \text{ is minimal}\} = \{a,c,d\}$. Hence we reach the next stage $\sigma_2 = (\langle\{a\},\{c\},\mathcal{W}\rangle, \langle\{a,b\},\{c\},\mathcal{W}\rangle, \langle\mathcal{W},\{a,c,d\},\mathcal{W}\rangle)$. Since we have now reached consistency, we end the process here with $f^{\mathcal{N}(R)}(\{a\},\{c\},\mathcal{W}) = \sigma_2$. From this we deduce $\mathbf{f}^{\mathcal{N}(R)}(\{a\},\{c\},\mathcal{W}) = \langle\mathcal{W},\{a,c,d\},\mathcal{W}\rangle$. For the corresponding merging operator we have $\Delta_{\mathcal{N}(R)}(\{a\},\{c\},\mathcal{W}) = \bigcap_{i=1}^{3} \mathbf{f}_i^{\mathcal{N}(R)}(\{a\},\{c\},\mathcal{W}) = \{a,c,d\}$.

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. In particular, this results in the following behaviour for the individual contraction functions $\ominus_i^{\mathcal{N}(R)}$:

Proposition 3 *Let R be a ranking assignment and let $i \in Sources^*$. Then, for all $S,T \in \mathcal{B}$, $S \ominus_i^{\mathcal{N}(R)} T = \{w \in \mathcal{W} \mid [R_i(S)](w) \leq [R_i(S)](T)\}$.*

From this the following can be shown:

Proposition 4 *Let R be a ranking assignment and let $i \in Sources^*$. Then the function $\ominus_i^{\mathcal{N}(R)}$ satisfies, in addition to* **(ic1)**–**(ic3)** *from Proposition 1, the following two properties:*

(ic4) $S \ominus_i^{\mathcal{N}(R)} (T_1 \cup T_2) \subseteq S \ominus_i^{\mathcal{N}(R)} T_1$

(ic5) *If* $(S \ominus_i^{\mathcal{N}(R)} (T_1 \cup T_2)) \cap T_1 \neq \emptyset$ *then* $S \ominus_i^{\mathcal{N}(R)} T_1 \subseteq S \ominus_i^{\mathcal{N}(R)} (T_1 \cup T_2)$

This means that $\ominus_i^{\mathcal{N}(R)}$ belongs to the class of contraction operators known as *severe withdrawal* operators, which were studied in (Rott & Pagnucco 1999). The rules **(ic4)** and **(ic5)** essentially correspond to the postulates $(\ddot{-}7\mathrm{a})$ and $(\ddot{-}8)$ given there. We also have the following nice characterisation of the merging operator $\Delta_{\mathcal{N}(R)}$:

Proposition 5 *Let R be a ranking assignment. Then, for all $\vec{S} \in \mathcal{B}^n$, $\Delta_{\mathcal{N}(R)}(\vec{S}) = \{w \in S_n \mid \max_{i \in Sources^*}[R_i(S_i)](w) \text{ is minimal}\}$.*

This "minimax" operator is a generalised version of the merging operator with integrity constraints Δ^{Max} given in (Konieczny & Pino-Pérez 1999), which employs a particular family of ranking assignments based on a notion of (symmetric) *distance* between propositional worlds. Similar operators are also discussed in (Meyer 2001, Meyer et al. 2001, Revesz 1993), and are shown to satisfy several interesting properties.

How do the SC functions $\mathbf{f}^{\mathcal{N}(R)}$ fare with regard to the minimal change postulates from Section 2.1? Well quite badly as it turns out. Indeed the ranking assignment R used in Example 2 provides a counter-example to show that the $\mathbf{f}^{\mathcal{N}(R)}$ do not, in general, satisfy even the weakest postulate **(sc5–)** mentioned there. To see this note that, in that example, we have $\{a\} \cap \mathbf{f}_2^{\mathcal{N}(R)}(\{a\},\{c\},\mathcal{W}) \cap \mathbf{f}_3^{\mathcal{N}(R)}(\{a\},\{c\},\mathcal{W}) = \{a\} \cap \{a,c,d\} \cap \mathcal{W} \neq \emptyset$. Now if $\mathbf{f}^{\mathcal{N}(R)}$ satisfied **(sc5–)** then we would conclude that $\mathbf{f}_1^{\mathcal{N}(R)}(\{a\},\{c\},\mathcal{W}) \subseteq \{a\} \ominus_1^{\mathcal{N}(R)} \overline{\{a\}}$. But $\mathbf{f}_1^{\mathcal{N}(R)}(\{a\},\{c\},\mathcal{W}) = \mathcal{W}$ and $\{a\} \ominus_1^{\mathcal{N}(R)} \overline{\{a\}} = \{a,b\}$. Hence $\mathbf{f}^{\mathcal{N}(R)}$ does not satisfy **(sc5–)**. Thus, interestingly, it seems that, while $\Delta_{\mathcal{N}(R)}$ might be quite well-behaved, there still seems to be room for improvement regarding the behaviour of $\mathbf{f}^{\mathcal{N}(R)}$.

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 $\mathbf{f}_n^{\mathcal{N}}(\vec{S}) = S_n$. So, using $\prec$ to denote a given strict total order on $Sources^*$ and assuming $i_1 \prec i_2 \prec \cdots \prec i_{n-1}$, we first focus on i_1 and repeatedly weaken S_{i_1} until it becomes consistent with S_n. The result of this weakening we will take to be $\mathbf{f}_{i_1}^{\mathcal{N}}(\vec{S})$.

Of course it may be that $S_{i_1} \cap S_n \neq \emptyset$ 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 $\mathbf{f}_{i_1}^{\mathcal{N}}(\vec{S}) \cap S_n$. 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.

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

$$g_2(\sigma) \quad = \quad \{i\}, \text{ where } i \in Sources^* \text{ is minimal under } \prec \text{ such that } S_i^m \cap (\bigcap_{j \prec i} S_j^m) \cap S_n^m = \emptyset$$

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 $\mathcal{N}(R, \prec) = \langle g_2, \{\blacktriangledown_\sigma\}_{\sigma\in\Sigma}\rangle$ where now g_2 is defined from $\prec$ 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 3.) Let's give a worked example of a belief negotiation model of this type.

Example 3 Suppose once more that $\mathcal{W} = \{a, b, c, d\}$, but this time that $Sources = \{1, 2, 3, 4\}$. We suppose that our sources provide the information profile $\vec{S} = (\{d\}, \{a, b, d\}, \{c\}, \{a, b, c\})$. We will use the belief negotiation model $\mathcal{N}(R, \prec)$, where $\prec$ is such that $1 \prec 2 \prec 3$ and the ranking assignment 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's construct the sequence $f^{\mathcal{N}(R,\prec)}(\vec{S}) = \sigma$ 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 $\prec$ and selecting the first one for which $S_i^0 \cap (\bigcap_{j\prec i} S_j^0) \cap S_4^0 = \emptyset$. Starting then with the minimal source in $Sources^*$, which is source 1, we immediately see that $S_1^0 \cap (\bigcap_{j\prec 1} S_j^0) \cap S_4^0 = S_1^0 \cap S_4^0 = \{d\} \cap \{a, b, c\} = \emptyset$. Hence source 1 is the loser of this negotiation round, i.e., $g_2(\sigma_0) = \{1\}$, and so must make some weakening. Since $\blacktriangledown_{\sigma_0}(1) = \{d\} \cup \{w \in \overline{\{d\}} \mid [R_1(\{d\})](w) \text{ is minimal}\} = \{a, b, d\}$ this leads us to the next stage $\sigma_1 = (\vec{S}^0, \vec{S}^1)$, where $\vec{S}^1 = (\{a, b, d\}, \{a, b, d\}, \{c\}, \{a, b, c\})$. 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 4, i.e., we have $S_1^2 \cap S_4^2 = \{a, b, d\} \cap \{a, b, c\} \neq \emptyset$. Hence source 1 weakens no further. According to our ordering $\prec$, we must consider source 2 next. But $S_2^2 \cap (\bigcap_{j\prec 2} S_j^2) \cap S_4^2 = S_2^2 \cap S_1^2 \cap S_4^2 = \{a, b, d\} \cap \{a, b, d\} \cap \{a, b, c\} \neq \emptyset$ 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 $\blacktriangledown_{\sigma_1}(3) = \{c\} \cup \{w \in \overline{\{c\}} \mid [R_3(\{c\})](w) \text{ is minimal}\} = \{c, d\}$ which leads us to the next stage $\sigma_2 = (\vec{S}^0, \vec{S}^1, \vec{S}^2)$ where $\vec{S}^2 = (\{a, b, d\}, \{a, b, d\}, \{c, d\}, \{a, b, c\})$. 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 $\blacktriangledown_{\sigma_2}(3) = \{c, d\} \cup \{w \in \overline{\{c, d\}} \mid [R_3(\{c\})](w) \text{ is minimal}\} = \{a, c, d\}$ leading to the next stage $\sigma_3 = (\vec{S}^0, \vec{S}^1, \vec{S}^2, \vec{S}^3)$ where now $\vec{S}^3 = (\{a, b, d\}, \{a, b, d\}, \{a, c, d\}, \{a, b, c\})$. This time we *have* reached consistency, so the process stops here with $f^{\mathcal{N}(R,\prec)}(\vec{S}) = \sigma_3$ and $\mathbf{f}^{\mathcal{N}(R,\prec)}(\vec{S}) = \vec{S}^3 = (\{a, b, d\}, \{a, b, d\}, \{a, c, d\}, \{a, b, c\})$. For the corresponding merging operator we get $\Delta_{\mathcal{N}(R,\prec)}(\vec{S}) = \bigcap_{i=1}^4 S_i^3 = \{a\}$.

Note that, by Proposition 2, the $\ominus_i^{\mathcal{N}(R,\prec)}$ are the same as the $\ominus_i^{\mathcal{N}(R)}$ from the previous subsection. Meanwhile we can characterise $\Delta_{\mathcal{N}(R,\prec)}$ with the help of the following piece of extra notation: We let $<_{lex}$ denote the lexicographic ordering on $\mathbb{N}^{n-1}$, i.e., given two tuples $\vec{x}, \vec{y} \in \mathbb{N}^{n-1}$ such that $\vec{x} = (x_1, \ldots, x_{n-1})$ and $\vec{y} = (y_1, \ldots, y_{n-1})$, 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 $i < j$. Then we have the following:

Proposition 6 *Let R be a ranking assignment and let $\prec$ be a strict total order on $Sources^*$. Then, assuming $i_1 \prec i_2 \prec \cdots \prec i_{n-1}$ and using r_j as an abbreviation for $[R_{i_j}(S_{i_j})]$, we have $\Delta_{\mathcal{N}(R,\prec)}(\vec{S}) = \{w \in S_n \mid (r_1(w), r_2(w), \ldots, r_{n-1}(w))$ is minimal under $<_{lex}\}$.*

Thus $\Delta_{\mathcal{N}(R,\prec)}(\vec{S})$ collects all the "best" worlds in S_n, in the special sense where one world is considered "better" than another if it is assigned lower rank by source i_1, or, in case they are assigned the same rank by i_1, it is assigned a lower rank by i_2, or, in case they are also assigned the same rank by i_2, it is assigned a lower rank by i_3, or, etc. Thus the effect when merging is that the opinion of source i is given *precedence* over that of i' whenever $i \prec i'$. Such a lexicographic approach to merging has been considered in (Meyer 2001) (see

Section 4.5 there) where the $\prec$ is interpreted as a given ordering of *reliability* on the sources, i.e., the most reliable sources are given precedence.

Finally, what can we say this time about the SC functions $\mathbf{f}^{\mathcal{N}(R,\prec)}$? First of all we may show the following:

Proposition 7 *Let R be a ranking assignment and $\prec$ a strict total order on Sources*. Then, for each information profile $\vec{S}$, we have*

$$\mathbf{f}_i^{\mathcal{N}(R,\prec)}(\vec{S}) = S_i \ominus_i^{\mathcal{N}(R,\prec)} \big(\mathbf{f}_n^{\mathcal{N}(R,\prec)}(\vec{S}) \cap \bigcap_{j \prec i} \mathbf{f}_j^{\mathcal{N}(R,\prec)}(\vec{S})\big).$$

In other words $\mathbf{f}_i^{\mathcal{N}(R,\prec)}(\vec{S})$ is equal to the result – according to i's individual contraction function relative to $\mathbf{f}^{\mathcal{N}(R,\prec)}$ – of weakening S_i to be jointly consistent with $\mathbf{f}_n^{\mathcal{N}(R,\prec)}(\vec{S})$ together with all the $\mathbf{f}_j^{\mathcal{N}(R,\prec)}(\vec{S})$ for which j precedes i according to $\prec$. Using this together with the fact that the $\ominus_i^{\mathcal{N}(R,\prec)}$ satisfy the properties **(ic4)** and **(ic5)** from Proposition 4 then allows us to prove:

Proposition 8 *Let R be a ranking assignment and $\prec$ a strict total order on Sources*. Then the SC function $\mathbf{f}^{\mathcal{N}(R,\prec)}$ satisfies* **(sc6)**.

Thus, imposing a strict "order of weakening" on the sources has forced our SC function to satisfy the "equilibrium" property **(sc6)** (and hence also **(sc5)** and **(sc5–)**).

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 (Booth 2001). 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 (Konieczny & Pino-Pérez 1999).

Acknowledgements

This work is supported by the DFG research project "Computationale Dialektik". Thanks are due to Thomas Meyer for a couple of helpful suggestions.

References

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

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) 137–150.

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

P. Gärdenfors, *Knowledge in Flux*, MIT Press (1988).

S. O. Hansson, *A Textbook of Belief Dynamics*, Kluwer Academic Publishers (1999).

S. Konieczny and 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) 488–498.

S. Konieczny and R. Pino-Pérez, Merging with integrity constraints, in *Proceedings of the Fifth European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty (ECSQARU '99)* (1999) 233–244.

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

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

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

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

T. A. Meyer, A. K. Ghose and 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) 466–477.

P. Z. 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) 71–82.

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

W. Spohn, Ordinal conditional functions: A dynamic theory of epistemic states, in W. L. Harper and B. Skyrms (eds.), *Causation in Decision, Belief Change and Statistics*, volume 2, (1988).

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

D. N. Walton and E. C. W. Krabbe, *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*, Albany, State University of New York Press (1995).

M. J. Wooldridge and S. Parsons, Languages for negotiation, in *Proceedings of the Fourteenth European Conference on Artificial Intelligence (ECAI 2000)* (2000) 393–400.