Interleaved Contractions

Wiebe van der Hoek[1] and Maarten de Rijke[2]

[1] Dept. of Computer Science, Utrecht University
P.O. Box 80089, 3508 TB Utrecht, the Netherlands
wiebe@cs.ruu.nl
[2] Dept. of Computer Science, University of Warwick
Coventry CV4 7AL, England
mdr@dcs.warwick.ac.uk

Abstract

We study an approach to concurrent contractions, that is, to simultaneous contractions performed by multiple agents. Using ideas from the semantics of programming we adopt an interleaved approach to reason about concurrent contractions. Although many of the notions from the traditional Gärdenfors approach transfer to this setting, our approach also forces us to depart from the Gärdenfors framework in important ways. We present laws describing rational concurrent contractions, as well as a construction that satisfies these laws.

1 Introduction

In real life concurrent accessing of data is the rule. Multiple agents are working on the same theory, and multiple copies of some data are kept in different locations. Typical examples include scientific research or writing a joint-publication, and practical applications vary from networks of personal computers and workstations sharing some common information to widely distributed applications such as automatic teller machines. The primary advantage of concurrent theory change as opposed to single agent theory change is the ability to share, access and engineer data in an efficient manner. The primary disadvantage is the added complexity required to ensure proper coordination between the agents taking part.

In a multi-agent setting, managing a belief set is a *concurrent* task: not only may several agents *retrieve* information from one and the same source, but it may also be the case that multiple agents have permission to *alter* a database (the flight booking procedures are a typical example here). What are sensible strategies for conflict resolution in case inconsistency strikes? The task of maintaining consistency in the setting of multi-agent theory change is more complex than in the single agent case, if only because of the many possibilities that become visible.

This paper is part of a larger project on concurrent theory change (see [7]). Its purpose is to demonstrate that concurrent theory change forms an interesting extension of the traditional Gärdenfors style approach towards theory change, one that has many faces and that calls for new tools. Here we will confine ourselves to the simplest case in which a number of agents have access to shared data. The data are changed via contractions, which may in principle be proposed by any one of the agents. We will explore some of the options and problems that present themselves. A central question of this paper is: assuming that multiple agents, each guided by a familiar set of rationality postulates, propose or perform (single agent) contractions for a shared theory, — what are the laws governing the global contractions?

The rest of the paper is organized as follows. In Section 2 we briefly outline the general set-up. Section 3 contains an informal discussion of concurrent contractions, and Section 4 recalls some facts from the standard Gärdenfors framework. Then, in Sections 5, 6 and 7 we present our formal approach to concurrent contractions, based on the idea of interleaving. We conclude the paper with comments and suggestions for further work in Section 8.

2 General Set-up

There have been many proposals to alter or extend the basic Alchourrón, Gärdenfors, Makinson (AGM) framework of theory change (see [5] for an overview), but most of the literature in the AGM tradition focuses on a single agent changing a theory as she receives new information. The actions of this solitary agent are usually specified in terms of functional input/output behavior:

$$(T, \phi) \mapsto T', \tag{1}$$

where the input consists of a collection of sentences T (the material to be changed) and a sentence ϕ (the newly received information), and the output is a collection of sentences T' (the result of the cognitive action). Traditionally, three forms of theory change are considered: *expansions*, where we add ϕ to T and close under logical consequence; *contractions*, where we remove ϕ from T while preserving as much of T as possible; and *revisions*, where we add ϕ to T while maintaining or restoring consistency. In this paper we change the format given in (1), and consider concurrent contractions that are specified by expressions of the form

$$T \sim \begin{pmatrix} \phi_1 \\ \vdots \\ \phi_n \end{pmatrix}, \tag{2}$$

or $T \sim \vec{\phi}$, where T is as before, and $\vec{\phi}$ is a vector of formulas to be contracted from T; $\sim$ is the concurrent contraction action whose principles we want to understand. The basic assumption here is that there are n agents $A_1, \ldots,$ A_n, each of whom proposes or performs a contraction of T in accordance with her own contraction operation. That is, A_1 proposes or performs a contraction of T by $\phi_1, \ldots, A_n$ proposes or performs a contraction of T by ϕ_n, where each agent A_i has her own contraction operation $-_i$. The expression in (2) denotes the result(s) of an operation on T that is somehow composed of contractions of T by $\phi_1, \ldots, \phi_n$ performed by, respectively, A_1, $\ldots, A_n$ using their respective contraction operations $-_1, \ldots, -_n$. The key questions we address are:

- How can we model concurrent contractions?
- Which laws govern the concurrent contraction operation $\sim$?
- How can $\sim$ be understood in terms of the single agent operations $-_1$, $\ldots, -_n$?

Below we will explore concurrent contractions. We leave the much more complicated (and realistic) case of *heterogeneous* concurrent theory change in which multi-agent contractions, revisions and expansions may take place concurrently to later publications.

3 Why Contract Concurrently?

Before proceeding we give an informal discussion of concurrent contractions. As outlined above, our basic picture is one where n agents $A_1, \ldots, A_n$ simultaneously want to remove information from a given background theory T, that is: each agent proposes or performs a contraction, using her private contraction operation.

To give an example of concurrent theory change at work, one can think of a patient's record in a medical database. Various agents contribute to the theory contained in the database: a family doctor's report, various laboratories with their test results, specialists with further information.... Clearly, it is important that consistency be preserved. One may conceptualize this is by personifying consistency checking in terms of a checker that performs consistency checks at certain discrete intervals. If the checker detects an inconsistency in the shared theory, she rings the alarm bell, asking the agents to suggest contractions that will help remove the inconsistency. The agents then perform or suggest a contraction. Having different areas of expertise, the agents are likely to base their suggested contractions on different notions of which information is more reliable (or 'epistemically entrenched') than other. In other words, when agents suggest a contraction for the shared theory they suggest both *which* information should be given up, and *how* this should be done in their opinion. Therefore, the global change that is to be made to the theory is in general composed out of a finite number of 'private' contractions being performed concurrently.

In the special case where all agents employ the same contraction function, there is a clear connection with the *multiple contractions* proposed by Fuhrmann and Hansson [3], and with forms of *iterated belief change* that have recently been described by Lehmann and others (see [10]).

Ideas related to concurrent contraction also appear in non-epistemic settings. For example, co-authoring and joint research are processes in which concurrent contractions occur frequently. They seem especially appropriate when bugs or inconsistencies are discovered in cases where agents have sole responsibilities for certain parts of the work, and each author can perform contractions on the parts for which she holds responsibility. And of course, in concurrent databases concurrent transactions occur all the time. It is difficult, however, to find pure cases of concurrent contractions that are substantially different from the above ones.

4 Laws and Models for Single Agent Contractions

In this section we describe the laws governing the contraction operations of individual agents taking part in a concurrent contraction; as explained above, we assume that each agent comes equipped with her own contraction function. We start with some technical preliminaries.

Our background language is simply propositional logic, equipped with a classical consequence operator Cn that satisfies all the usual properties (see [4]). A *theory* is a set of formulas T that is closed under Cn; a *belief base* K is a set of formulas that needs not be a theory. In the AGM tradition there are two ways of reasoning about contraction functions, a *syntactic* way which specifies postulates that reasonable contraction functions should satisfy, and a *semantic* way that defines contractions functions obeying those laws. Here's a list of the standard AGM postulates for contraction.

$T - \phi$ is a theory (logically closed) whenever T is	(Closure)
$T - \phi \subseteq T$	(Inclusion)
If $\phi \notin T$, then $T - \phi = T$	(Vacuity)
If $\nvdash \phi$, then $\phi \notin T - \phi$	(Success)
If $\phi \in T$, then $T \subseteq \mathrm{Cn}((T - \phi) \cup \phi)$	(Recovery)
If $\vdash \phi \leftrightarrow \psi$ then $T - \phi = T - \psi$	(Extensionality)

We refer the reader to [4, 5] for a discussion. The above laws constrain how contraction functions $-$ should operate on a single, fixed theory T. But when n agents each come up with a formula ϕ_i to be contracted from a theory T, they should not only provide the system with a contraction function $-_i$, but, since the actual implementation of $T \sim \vec{\phi}$ may deal with several 'intermediate' results T' from which some of the ϕ_i's still have to be contracted, their contraction functions should indicate how to remove ϕ_i from arbitrary theories.

Hansson [6] gives a formal account of contraction functions able to deal with arbitrary theories. His approach is formulated in terms of belief bases K rather than theories T, and he moreover allows for contractions with *sets* of formulas rather than single formulas. We reformulate Hansson's original postulates for the 'base/set' case for the 'theory/formula' case.

Definition 1 (Postulates for single agent contraction) We propose the following postulates for a single agent contraction function $-$ that is defined for any theory T and formula ϕ:

$T - \phi$ is a theory (logically closed) whenever T is (Closure)

$T - \phi \subseteq T$ (Inclusion)

If $\psi \in T \setminus (T - \phi)$ then there exists T' with $T - \phi \subseteq T' \subseteq T$ and $T' \nvdash \phi$, but $T', \psi \vdash \phi$ (Relevance)

If $T' \vdash \phi \leftrightarrow \psi$ for all subtheories $T' \subseteq T$, then $T - \phi = T - \psi$ (Uniformity)

If $\nvdash \phi$, then $\phi \notin T - \phi$ (Success)

Relevance ensures that if a formula ψ is excluded from T when ϕ is rejected, then ψ plays a role in the fact that T implies ϕ. Whereas Success ensures that formulas that should be given up are in fact given up, Relevance blocks the deletion of formulas that need not be deleted. Uniformity ensures that the result of contracting T with ϕ depends only on the subsets of T that imply ϕ; if all subsets derive a given formula ϕ iff they derive ψ, then contracting with either ϕ or ψ produces the same result. Observe that Vacuity is derivable from Inclusion and Relevance.

In the setting of concurrent contractions it may well be that some agents want to refrain from action. The next proposition shows how we can mimic this situation.

Proposition 2 *If a contraction function $-$ satisfies the postulates of Definition 1, then, for any theory T and tautology $\top$, we have $T - \top = T$.*

The best known model of a contraction function in the AGM theory is partial meet contraction. It is defined as follows. Let $T \perp \phi$ denote the set of maximal subsets of T that fail to imply ϕ. A *one-place selection function* for T is a function s such that for all formulas ϕ, if $T \perp \phi$ is non-empty, then $s(T \perp \phi)$ is a non-empty subset of $T \perp \phi$. When $T \perp \phi$ is empty, $s(T \perp \phi) = \{T\}$. Then, an operation $-$ on a theory T is a *partial meet contraction* if $T - \phi$ is the intersection of the selected maximal subsets of T that fail to imply ϕ: $T - \phi = \bigcap s(T \perp \phi)$.

One-place selection functions are specific for a particular theory; if s is a one-place selection function for T, and $T \neq T'$, then s is not a one-place selection function for T' (see Hansson [6]). Selection functions that work for arbitrary theories are obtained by extending them with an additional argument; thus we will assume that each agent i is equipped with a two-placed selection function s, where, for each theory T and set of theories $(S \perp \psi)$, we have $s(T, (S \perp \psi)) \subseteq (S \perp \psi)$.

The following result links up the postulates for $-$ with two-placed contraction functions; a proof is given in the Appendix.

Theorem 3 *A single agent contraction function $-$ satisfies the postulates of Definition 1 iff there exists a two-placed selection function s with $T - \phi = \bigcap s(T, (T \perp \phi))$, for any theory T and formula ϕ.*

Now that we have shown how an agent's contractions can be modeled using two-placed contraction functions s, we pause a moment and reflect upon the desired effects of the first argument of s. Recall that $S \perp \psi$ denotes all maximal sub-theories of S that do not entail ψ (if $\not\vdash \psi$). When contracting ψ from S, the function s should make a selection from these sub-theories. This selection should principally reflect the agent's preferences among the theories in $(S \perp \psi)$. Thus, if we have

$$(S \perp \psi) = (U \perp \chi) \neq \emptyset,$$

it seems natural to require that

$$\bigcap s(S, (S \perp \psi)) = \bigcap s(U, (U \perp \chi)).$$

In other words, the common parts of the selections agree whenever possible.[1] Hansson calls a selection function with this property *unified*. When working with belief bases this property doesn't come for free. Hansson comes up with a condition on contraction functions called *redundancy* to characterize unified partial meet contractions. In our set up this redundancy principle reads as follows:

Redundancy reformulated. Suppose T is a theory, and $\not\vdash \phi$. Suppose furthermore that Z is a set of formulas, satisfying: (i) $T \cup Z$ is a theory, and (ii) for all $\zeta \in Z$: $\vdash \zeta \to \phi$. Then we have: $T - \phi = (T \cup Z) - \phi$.

Theorem 4 *If a contraction function $-$ satisfies the postulates of Definition 1, then it also satisfies redundancy.*

[1]Note that the first argument of s is still relevant: when modeling a contraction $T - \phi$, we calculate $s(T, (T \perp \phi))$.

Theorem 4 (the proof of which is to be found in the Appendix) guarantees that we do not have to add redundancy as a separate postulate, so that we can now formulate the main result of this Section; its proof is given in the Appendix.

Theorem 5 *A single agent contraction function $-$ satisfies the postulates of Definition 1 iff there exists a two-placed* unified *selection function s with $T - \phi = \bigcap s(T, (T \perp \phi))$, for any theory T and formula ϕ.*

In the sequel, we will assume that selection functions are unified, and we will often suppress their first argument.

5 From Sequential to Interleaved Contractions

In many models of situations in which multiple agents need to access shared resources, one finds a reduction to a sequential, non-deterministic scheme. Our model of concurrent contractions will be based on the same idea. To see how we arrive at our model, consider the following diagram in which a contraction by a single agent i is depicted by a line segment labeled with i.

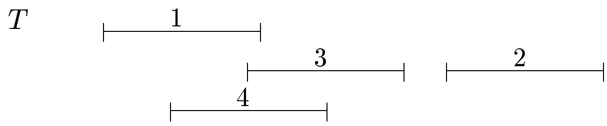

Figure 1: Overlapping contractions

It pictures how multiple agents might — in principle — act on a single theory T to perform their individual contractions as time progresses: their actions might or might not overlap in arbitrary ways. But what does it mean for an agent i to start a contraction *while another agent j is still performing her contraction*? To what should i apply her selection function? What should she act on, if not on the outcome of j's actions? To perform a concurrent contraction one should execute the individual single agent contractions, one at a time. Thus, instead of Figure 1, Figure 2 seems to offer a much more realistic picture.

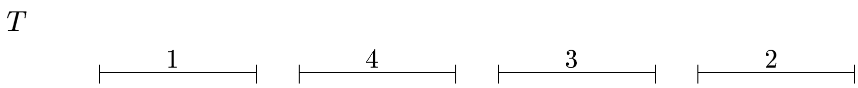

Figure 2: Interleaved contractions

To understand this situation, it may help to observe that there is a clear analogy with some forms of concurrent computation. Specifically, the situation is reminiscent of the concurrent execution of several independent programs on a single processor (see e.g. [2]). In a popular formal model *concurrency* is represented by *interleaving*. This means that parallel processes are never executed at precisely the same instant, but take turns in executing atomic transitions. When one of the participating processes executes an atomic transition, the others are inactive. Thus, rather than input/output pairs, execution *sequences* of the atomic instructions of sequential processes are at the focus of attention. And rather than talk about input/output pairs, one describes properties of concurrent programs that hold under some or all interleavings of the instructions. Let us briefly expand on this issue.

As parallel execution of sequential processes is modeled by the non-deterministic interleaving of atomic steps of the individual processes in interleaved models of concurrent programs, a program starting in a given state may follow any one of a number of computation paths corresponding to the different non-deterministic choices the program might make. The different computation paths thus represent alternative possible 'futures': at each moment, time may split into alternative courses and thus has a 'branching' tree-like structure. A semantic theory of computations provides a formal basis for describing or deducing properties of programs under *all* possible interleavings (see [2] for further details).

A similar concern is found in concurrent database theory, where one studies mechanisms for controlling the execution of several transactions at the same time. Here, one of the main interests lies in describing all possible executions of transactions and in identifying serializable transactions, that is: transactions that are equivalent in some sense to serial (consistency preserving) database transactions (see [9, Chapter 10] for an introductory overview).

In our setting of multi-agent contractions, we take a similar interleaved approach. Concurrent contractions will be viewed as (collections of) sequences of 'atomic' single agent contractions that don't overlap and that don't interfere. This interleaved approach calls for new ways of thinking about theory change. For a start, if we reduce concurrent contractions to non-deterministic sequential contractions, instead of single one step contractions we should be considering collections of sequences of contractions that are organized in a tree as in Figure 3. But then, we also have to give up the idea of concurrent contractions as functions. For, in general, even singleton contractions are not commutative, that is: the identity

$$(T - \phi_1) - \phi_2 = (T - \phi_2) - \phi_1$$

is not universally valid (see [6] for a plausible counterexample). Hence, even in the case where we only have two agents that share the same contraction operation, say $-_1$, the global concurrent contraction $T \sim \binom{\phi_1}{\phi_2}$ may have two interleaving executions leading to different results: $(T -_1 \phi_1) -_1 \phi_2$ and

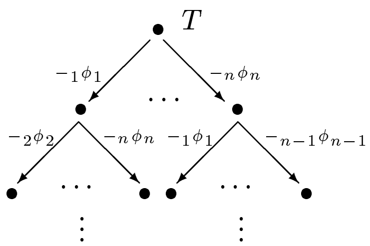

Figure 3: Interleaved contractions

$(T -_1 \phi_2) -_1 \phi_1$. As a consequence, it makes little sense to talk about *the* outcome of interleaved contractions. As a further consequence, we have to abandon the idea that contractions can be specified in terms of preconditions and postconditions. Instead, we need to reason about intermediate stages of a concurrent contraction, as these clearly have internal structure in the interleaved approach.

6 Laws for Interleaved Contractions

By and large, theory change in the AGM tradition has had two main concerns: (1) Constraints and axioms that rational operations of theory change should satisfy; and (2) Explicit constructions of operations of theory change that satisfy those constraints or axioms. In our approach to concurrent theory change we follow the same strategy. In particular, in this section we present a list of postulates that any reasonable operation of interleaved contraction should satisfy. Then, in Section 7 below, we present a model for interleaved contractions that satisfies these postulates.

We first need some notation. Let $\vec{\phi}$ and $\vec{\psi}$ be vectors of formulas of length n. We write $\vec{\phi} \equiv \vec{\psi}$ for 'for all $i \leq n$: $\vdash \phi_i \leftrightarrow \psi_i$,' and for a set of formulas Z, $\vec{\phi} \equiv_Z \vec{\psi}$ means that for all $i \leq n$: $Z \vdash \phi_i \leftrightarrow \psi_i$. Special vectors are $\vec{\top}$ and $\vec{\bot}$, consisting only of the formulas $\top$ and $\bot$, respectively; $\vec{\phi}[\chi/\phi_i]$ denotes the result of replacing the i-th component ϕ_i by χ. A *concurrent contraction function* is a function that produces a set of theories when given a theory T and a vector $\vec{\phi}$ as input. Finally, we write $(T \sim \vec{\phi}) \propto S$ for 'S is a result of concurrently contracting T with $\vec{\phi}$.'

Definition 6 (Basic postulates for interleaved contractions) Let T be a set of formulas, and let $n > 1$ be a natural number (the number of agents taking part). We assume that each $-_i$ satisfies the postulates for single agent contraction given in Definition 1.

If T is a theory and $(T \sim \vec{\phi}) \propto S$, then S is also a theory (C-closure)

If $(T \sim \vec{\phi}) \propto S$ then $S \subseteq T$ (C-inclusion)

Suppose $(T \sim \vec{\phi}) \propto S$ and $\psi \notin S$, $\psi \in T$. Then there are T' and $i \leq n$ such that $S \subseteq T' \subseteq T$, $T' \not\vdash \phi_i$, and $T', \psi \vdash \phi_i$ (C-relevance)

If $\vec{\phi} \equiv_{T'} \vec{\psi}$ for all subtheories $T' \subseteq T$, then $(T \sim \vec{\phi}) \propto S$ iff $(T \sim \vec{\psi}) \propto S$ (C-uniformity)

For all $i \leq n$, if $\not\vdash \phi_i$ and $(T \sim \vec{\phi}) \propto S$, then $\phi_i \notin S$ (C-success)

If for all $j \neq i$, we have $\phi_j = \top$, then $(T \sim \vec{\phi}) \propto S$ iff $S = T -_i \phi_i$ (Solo)

C-closure and C-inclusion are straightforward concurrent versions of their single agent counterparts. C-Relevance says that for every formula ψ that is given up in a concurrent contraction from T, there is an agent that is 'responsible' for this removal; according to the individual obedience to Relevance, this agent can determine a part of T from which the formula he proposed for contraction has indeed been removed, and in the process of doing this, ψ had to be given up. C-uniformity says that if no subtheory T' of T can distinguish any component ϕ_i of $\vec{\phi}$ from the corresponding component ψ_i of $\vec{\psi}$, then concurrently contracting $\vec{\phi}$ from T cannot be distinguished from concurrently contracting $\vec{\psi}$ from T. This uniformity postulate implies the following condition of C-extensionality:

If $\vec{\phi} \equiv \vec{\psi}$ then $(T \sim \vec{\phi}) \propto S$ iff $= (T \sim \vec{\psi}) \propto S$ (C-extensionality)

It guarantees that only the *content* of the individual's proposal for contraction matters, not the actual *form*. C-success guarantees that, as long as an agent does not propose to contract a tautology, her request for contraction will have been granted in each of the possible results. Thus, whereas C-relevance says that each formula that is given up in a concurrent contraction should be due to one of the agents, C-success guarantees that all of the agents' wishes will be met as far as they are reasonable. Finally, the Solo postulate shows that interleaved contractions really builds on the individual contraction strategies: when only one agent comes up with a non-trivial formula to be removed, it will be her strategy that determines the result of the concurrent contraction.

The postulates in Definition 6 provide no means to reason about possible 'intermediate' results of interleaved contractions, and they certainly don't impose the condition that the concurrent contraction process can be unraveled into successive single agent contractions. To make up for this, we consider two further laws: Decomposition and Composition.

(Decomposition)

$$(T \sim \vec{\phi}) \propto S \Rightarrow \begin{cases} \textit{either } T = S \text{ and } \vec{\phi} \equiv \vec{\top} \\ \textit{or} \text{ there exist } S' \text{ and } i \text{ with } i \leq n \text{ and} \\ S \subseteq S' \subseteq T \text{ such that } \phi_i \not\equiv \top,\ S' = T -_i \phi_i \\ \text{and } (S' \sim \vec{\psi}) \propto S, \text{ where } \vec{\psi} = \vec{\phi}[\top/\phi_i] \end{cases}$$

Decomposition says that concurrently contracting with $\vec{\top}$ is a void action and that a concurrent contraction with $\vec{\phi} \neq \vec{\top}$ can be decomposed in an individual contraction $-_i$ followed by another, yet simpler, concurrent contraction.

(Composition)

$$(T \sim \vec{\phi}) \propto S \Leftarrow \begin{cases} \textit{either } T = S \text{ and } \vec{\phi} \equiv \vec{\top} \\ \textit{or} \text{ there exist } S' \text{ and } i \text{ with } i \leq n \text{ and} \\ S \subseteq S' \subseteq T \text{ such that } \phi_i \not\equiv \top,\ S' = T -_i \phi_i \\ \text{and } (S' \sim \vec{\psi}) \propto S, \text{ where } \vec{\psi} = \vec{\phi}[\top/\phi_i] \end{cases}$$

Composition states that if one recursively unravels a concurrent contraction $T \sim \vec{\phi}$ into an individual contraction $T -_i \phi_i$ followed by a concurrent contraction of a vector $\vec{\psi}$ (obtained from $\vec{\phi}$) from the theory $(T -_i \phi_i)$, one ends up with a theory S that will be a result of the initial concurrent contraction. Notice that the Solo postulate is a consequence of Decomposition. If we think about interleaving contractions in an algorithmic way, we can view the Composition and Decomposition postulates as halting criteria: to contract $\vec{\phi}$ from T, try to turn all components of $\vec{\phi}$ into the formula $\top$ by successively contracting with one ϕ_i after another until $\vec{\phi}$ equals $\vec{\top}$.

Observe that the conjunction of Decomposition and Composition postulates is equivalent to the following statement; let n be the length of $\vec{\phi}$.

$$(T \sim \vec{\phi}) \propto S \quad \text{iff} \quad \text{there exists a permutation } f \text{ of } \{1, \ldots, n\} \text{ such that } S = ((\cdots (T -_{f(1)} \phi_{f(1)}) \cdots) -_{f(n)} \phi_{f(n)}).$$

Theorem 7 *Assume that a set of individual contraction functions $-_i$ ($1 \leq i \leq n$) and a concurrent contraction $\sim$ are connected via the Decomposition and Composition laws. If all the $-_i$'s satisfy the postulates from Definition 1, then $\sim$ satisfies all the Concurrent postulates from Definition 6.*

Proof. Suppose all the $-_i$'s satisfy the postulates from Definition 1. As pointed out above, we have that $(T \sim \vec{\phi}) \propto S$ iff for some permutation f of $\{1, \ldots, n\}$

$$S = ((\cdots (T -_{f(1)} \phi_{f(1)}) \cdots) -_{f(n)} \phi_{f(n)}).$$

Let $T_0 = T$ and $T_i = T_{i-1} -_{f(i)} \phi_{f(i)}$, for $i > 0$. Note that, by Inclusion, we have $T_i \subseteq T_{i-1}$ (for $1 \leq i \leq n$). Now, $\sim$ satisfies C-closure trivially: if T is a theory then, by n applications of Closure, $T_0, T_1 \ldots T_n = S$ are all theories. C-inclusion follows similarly.

For C-relevance, suppose all $-_i$'s satisfy Relevance, and suppose that $(T \sim \vec{\phi}) \propto S, \phi \notin S, \psi \in T$. Since each $-_i$ satisfies Inclusion, there must be some j such that $\psi \in T_j$, $\psi \notin T_{j+1}$. Since $T_{j+1} = T_j -_{f(j)} \phi_{f(j)}$ and $-_{f(j)}$ satisfies relevance, we find a T' with $T_j -_{f(j)} \phi_{f(j)} \subseteq T' \subseteq T_{j+1}$, $T' \not\vdash \phi_{f(j)}$

and $T', \psi \vdash \phi_{f(j)}$. Using Inclusion, we see that $S = T_n \subseteq T' \subseteq T_0 = T$. From this we can conclude that $\sim$ satisfies C-relevance.

For C-uniformity, suppose that $\vec{\phi} \equiv_{T'} \vec{\psi}$ for all subtheories $T' \subseteq T$. Thus, for all $i \leq n$, $T' \vdash \phi_i \leftrightarrow \psi_i$. Suppose furthermore that $(T \sim \vec{\phi}) \propto S$: we have to demonstrate that $(T \sim \vec{\psi}) \propto S$. But, since all the $-_i$'s satisfy Relevance, we immediately see that

$$S = (T -_{f(1)} \phi_{f(1)}) \cdots) -_{f(n)} \phi_{f(n)}) = (T -_{f(1)} \psi_{f(1)}) \cdots) -_{f(n)} \psi_{f(n)}),$$

which proves that $(T \sim \vec{\psi}) \propto S$.

For C-success, suppose $\not\vdash \phi_i$ and $(T \sim \vec{\phi}) \propto S$. Let $i = f(k)$, then, by Success, $\phi_i \notin T_k = T_{k-1} -_i \phi_i$ and, by Inclusion, $\phi \notin T_n = S$.

Finally, we prove that $\sim$ satisfies Solo: suppose that for all $j \neq i$, we have $\phi_j = \top$. Let k be such that $i = f(k)$. Then, by Proposition 2, we have for any $m \neq k$, that $T_m = T_{m_1}, m > 0$. Thus, we have

$$T = T_0 = T_1 = \ldots = T_{k_1}, T_k = T_{k_1} -_i \phi_i = T_{k+1} = \ldots = T_n = S.$$

Thus, $S = T -_i \phi_i$. $\dashv$

Theorem 7 expresses a *transfer* property: if we define a concurrent contraction $\sim$ via Composition and Decomposition using individual contractions $-_i$, we get the rationality postulates for $\sim$ if we impose rationality postulates on all the $-_i$'s. Theorem 8 expresses a *projection* principle going in the converse direction.

Theorem 8 *Assume that a set of individual contraction functions $-_i$ ($1 \leq i \leq n$) and a concurrent contraction $\sim$ are connected via the Decomposition and Composition laws. If $\sim$ satisfies the Concurrent postulates from Definition 6, then all the $-_i$'s satisfy the postulates from Definition 1.*

Proof. We note the following. For any formula ϕ, let $\vec{v}(i, \phi)$ be the vector with ϕ at index i, and with $\top$ at all other places: $\vec{v}(i, \phi)_i = \phi$ and $\vec{v}(i, \phi)_j = \top, i \neq j$. Using the interleaved contraction postulate Solo, we immediately obtain:

$$(T \sim \vec{v}(i, \phi)) \propto S \Leftrightarrow S = T -_i \phi \qquad (3)$$

Equation (3) expresses that a single agent contraction can be modeled by the multiple-agent contraction, provided that all agents but one refrain from acting. Now, let $\sim$ satisfy the properties of Definition 6. Then, using (3), one easily reads off the $-_i$ properties Closure, Inclusion and Success from C-closure, C-inclusion and C-success for $\sim$, respectively. For Relevance, suppose that $\psi \in T \setminus (T -_i \phi)$. Using C-relevance, we find a $j \leq n$ and a T' with $(T -_i \vec{v}(i, \phi)) = S \subseteq T' \subseteq T$ such that $T' \not\vdash \vec{v}(i, \phi)_j$ and $T', \psi \vdash \vec{v}(i, \phi)_j$. Since for all $k \neq i$, $\vec{v}(i, \phi)_k = \top$, we must have $j = i$. Since $\vec{v}(i, \phi)_i = \phi$, we have $T' \not\vdash \phi$ and $T', \psi \vdash \phi$ for some T' with $(T -_i \phi) \subseteq T' \subseteq T$, expressing that $-_i$ satisfies Relevance. To check Uniformity for $-_i$, suppose that $T' \vdash \phi \leftrightarrow \psi$ for all subtheories $T' \subseteq T$. By definition of $\vec{v}(i, \phi)$, we immediately see that $\vec{v}(i, \phi) \equiv_{T'} \vec{v}(i, \psi)$ so that C-uniformity yields

$$(T \sim \vec{v}(i, \phi)) \propto S \text{ iff } (T \sim \vec{v}(i, \psi)) \propto S.$$

Using (3) we conclude that $(T -_i \phi) = S = (T -_i \psi)$, and this proves Uniformity. $\dashv$

Combining Theorems 7 and 8 we see that the Composition and Decomposition postulates properly link individual and concurrent contractions together: postulates for individual operations are guaranteed by imposing postulates on the concurrent one, and vice versa.

Requiring that single agent contractions $-_i$ and a concurrent contraction $\sim$ are related through Composition and Decomposition is a non-trivial requirement, even if the single agent contractions satisfy the postulates of Definition 1, and the concurrent contraction satisfies the postulates of Definition 6. The main reason is that the postulates for contraction don't, in general, uniquely pin down its actual implementation. One can have different single contractions $-_1$ and $-_{1'}$ satisfying the postulates of Definition 1 (for example, $-_1$ can be a full meet contraction $-_1$, and $-_{1'}$ a partial meet contraction). Now, assume that $-_1$ and $-_{1'}$ are composed with $-_2, \ldots, -_n$ (all satisfying the postulates of Definition 1) into $\sim$ and $\sim'$, respectively, using Composition. Then, the single agent contractions $-_1, \ldots, -_n$ and the concurrent contraction $\sim'$ are not related via Composition and Decomposition.

7 A Model for Interleaved Contractions

Let $n > 1$ be the number of agents. We assume that for each i, agent i's contraction function is defined using a selection function s_i, as outlined in Section 4. The models we are about to define are called selection systems; they are based on the selection functions contributed by the individual agents. Roughly, a selection system is a collection of compositions of single agent selection functions that satisfies certain constraints.

More precisely, a *selection system* $(\mathcal{S}, \mathcal{T}, \mathsf{s}_0)$, intended to represent interleaved contractions, is given by the following components:

- $\mathcal{S}$, a set of states. Each state s is labeled with a theory $\mathrm{Th}(\mathsf{s})$. These are the theories that the theory of the *initial state* $\mathrm{Th}(\mathsf{s}_0)$ can evolve into by applying sequences of single agent contractions. Two states may be labeled with the same theory.
- $\mathcal{T}$, a set of possible transitions built up from the individual agents' single contraction:

$$\mathcal{T} = \{(\mathsf{s}, \mathsf{s}') \mid \exists i \exists \phi \, (\bigcap_i s_i(\mathrm{Th}(\mathsf{s}), (\mathrm{Th}(\mathsf{s}) \perp \phi)) = \mathrm{Th}(\mathsf{s}'))\}.$$

Here s' is called a *successor* (or $-_i\phi$*-successor*) of s; notation: $\mathsf{s} \xrightarrow{i,\phi} \mathsf{s}'$. For technical reasons we will assume that all successor steps are irreflexive: if $\mathsf{s} \xrightarrow{i,\phi} \mathsf{s}'$ then $\mathsf{s} \neq \mathsf{s}'$.

A state $\mathsf{s} \in \mathcal{S}$ is *terminal* if it has no successors. A *choice sequence* of a selection system $(\mathcal{S}, \mathcal{T}, \mathsf{s}_0)$ is a finite sequence $\sigma : \mathsf{s}_1, \ldots, \mathsf{s}_m$ satisfying the following requirements. First, the *Initiation* requirement says the state s_1 is the initial state of the selection system, that is: $\mathsf{s}_1 = \mathsf{s}_0$. Second, the *Consecution* requirement says that for each pair of consecutive states s_j, $\mathsf{s}_{j+1} \in \sigma$ there is a selection function s_i and a formula ϕ such that $\mathsf{s}_j \xrightarrow{i,\phi} \mathsf{s}_{j+1}$. (Observe that two states may be connected by multiple transitions.) Finally, the *Termination* requirement says that the final state s_m is a terminal state.

A *prefix* is a sequence $\mathsf{s}_1, \ldots, \mathsf{s}_k$ satisfying the requirements of initiation and consecution, but not necessarily of termination. The *length* of a prefix is its number of states.

Let T be a theory, n the number of agents, and $\vec{\phi} = (\phi_1, \ldots, \phi_n)$ a sequence of formulas. Our next aim is to determine what it means for a selection system $\mathbf{T} = (\mathcal{S}, \mathcal{T}, \mathsf{s}_0)$ to model or represent the interleaved contraction $T \sim \vec{\phi}$. We will impose three constraints. First, the *Start* constraint says that $\mathrm{Th}(\mathsf{s}_0)$, the theory of the initial state, should equal T. Second, the *Tightness* constraint requires that for every choice sequence σ in $\mathbf{T}$ and every $i \leq n$ there exists at most one pair of consecutive states s_j, s_{j+1} in σ such that $\mathsf{s}_j \xrightarrow{i,\phi} \mathsf{s}_{j+1}$. Intuitively, the tightness property says that no attempt is made to carry out a single agent contraction in $T \sim \vec{\phi}$ twice. Third, the *Fairness* constraint says that for every choice sequence σ in $\mathbf{T}$ and every $i \leq n$ there is a consecutive pair s_j, s_{j+1} in σ such that $\mathsf{s}_j \xrightarrow{i,\phi} \mathsf{s}_{j+1}$. holds. The fairness property expresses that every single agent contraction in $T \sim \vec{\phi}$ will eventually be carried out.

Let $\mathbf{S} = (\mathcal{S}, \mathcal{T}, \mathsf{s}_0)$ be a selection system. $\mathbf{S}$ is called a *model* for $T \sim \vec{\phi}$ if it satisfies the starting, tightness and fairness conditions for $T \sim \vec{\phi}$. Given a model $\mathbf{S} = (\mathcal{S}, \mathcal{T}, \mathsf{s}_0)$ for $T \sim \vec{\phi}$, a *proper choice sequence* of $T \sim \vec{\phi}$ is simply a choice sequence in $\mathbf{S}$. What this definition boils down to is that we view interleaved contractions as generators of proper choice sequences.

To be able to express the connection between concurrent contraction functions and selection systems, we say that a contraction function $\sim$ *generates a full selection system for T and $\vec{\phi}$* if there are single agent selection functions $s_1, \ldots, s_n$ such that $-_j$ is defined in terms of s_j $(1 \leq j \leq n)$, and for all S such that $(T \sim \vec{\phi}) \propto S$ there exists a sequence $S_0, \ldots, S_n$ such that $S_0 = T$, $S_{i+1} = \bigcap s_{f(i+1)}(S_i, (S_i \perp \phi_{f(i)}))$, where f is a permutation of $\{1, \ldots, n\}$, and $S_n = S$.

Theorem 9 *Let $\sim$ be a concurrent contraction function, and $-_i$ a set of single agent contractions. Then $\sim$ satisfies the C-postulates from Definition 6, and $\sim$ and $-_i$ are related via the Composition and Decomposition* *laws from Section 6 iff, for every theory T and vector of formulas $\vec{\phi}$, $\sim$ generates a full selection system for T and $\vec{\phi}$.*

Proof. First, suppose that $\sim$ satisfies the C-postulates of Definition 6, and suppose also that $\sim$ and $-_i$ are related via Composition and Decomposition. Let S be such that $(T \sim \vec{\phi}) \propto S$. Just as in the proof of Theorem 7 we find a permutation f of $\{1, \ldots, n\}$ such that

$$S = ((\cdots (T -_{f(1)} \phi_{f(1)}) \cdots) -_{f(n)} \phi_{f(n)}).$$

Now, define $S_0 = T$ and $S_{i+1} = (S_i -_{f(i+1)} \phi_{f(i+1)})$. Theorem 8 guarantees that each individual contraction $-_{f(i+1)}$ satisfies the postulates of Definition 1 and hence, we may use one direction of Theorem 3 to conclude that each contraction $S_i -_{f(i+1)} \phi_{f(i+1)}$ corresponds to taking the intersection of the selection that agent $f(i)$ generates, using S_i and $S_i \perp \phi_{f(i)}$, so that we have $S_{i+1} = \bigcap s_{f(i+1)}(S_i, (S_i \perp \phi_{f(i)}))$. This proves that every T and $\vec{\phi}$ generate a full selection system.

For the converse, suppose that for every T and $\vec{\phi}$, the operator $\sim$ generates a full selection function. Let S, T and $\vec{\phi}$ be such that $(T \sim \vec{\phi}) \propto S$. We know that, semantically, this gives rise to a sequence $S_0 = T$ and $S_{i+1} = \bigcap s_{f(i+1)}(S_i, (S_i \perp \phi_{f(i)}))$, where each $s_{f(i+1)}$ is a selection function. Now, we use the other direction of Theorem 3 to lift this semantic result to a syntactic level: we can associate a single agent contraction $-_{f(i+1)}$ satisfying the postulates of Definition 1 with each selection function $s_{f(i+1)}$, and we may write $S_{i+1} = (S_i -_{f(i+1)} \phi_{f(i+1)})$. Hence, by an application of Theorem 7 we conclude that $\sim$ satisfies the C-postulates of Definition 6. Finally, by observing that each sequence $S_0, \ldots, S_n$ in a full selection system determines a permutation f of $\{1, \ldots, n\}$ such that $S_{i+1} = S_i -_{f(i+1)} \phi_{f(i+1)}$, it follows that $\sim$ and $-_i$ are related via the Composition and Decomposition postulates. $\dashv$

With the above result we can 'complete the square' in the following diagram:

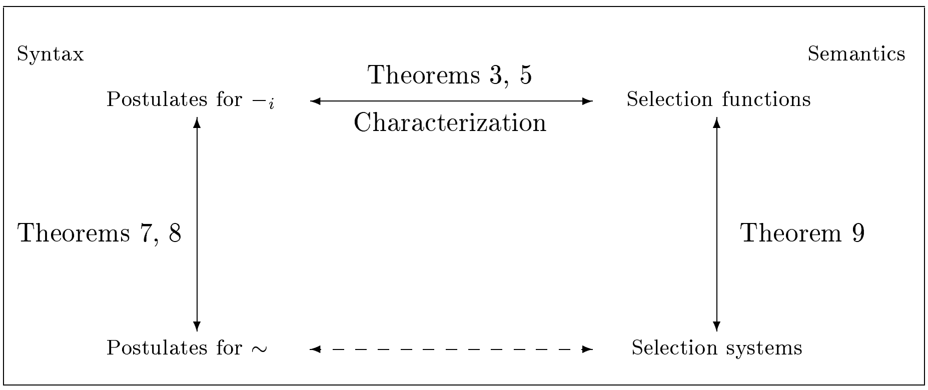

By walking around the above diagram we see that full selection systems are a model for our postulates for interleaved contraction, and any full selection system for T and $\vec{\phi}$ is given by the postulates for $\sim$.

8 Concluding Remarks

We have shown that concurrent contractions are well-behaved in that they satisfy a set of fairly transparent rationality postulates. On the assumption that all the underlying single agents contract in a rational way, and that concurrency is modeled in an interleaving manner.

In the course of the paper we have had to make explicit and alter some of the assumptions underlying the AGM approach to theory change as they seem no longer appropriate in our setting:

- In our interleaved setting theory change operations need not be functional; they are always defined but they need not have a unique outcome. (A similar deviation from the AGM assumptions in the context of single agent theory change is explored by Lindström and Rabinowitz [11].)

- In our interleaved setting theory change no longer is a one step operation. Although theory change typically occurs in dynamic environments in which agents may learn new information in a continuous process, the traditional AGM framework consistently avoids mentioning iterations of its operations. Recently a number of authors have abandoned this assumption, and considered forms of iterated theory change; see for example Lehmann [10] and Kfir-Dahav and Tennenholtz [8].

- In our interleaved setting theory change operations have internal structure, and they are no longer fully characterized by their pre-conditions and postconditions. In contrast, the traditional AGM postulates have nothing to say about the internal mechanisms by which operations of theory change achieve their goals.

To formulate it in a single sentence, the reason that the above assumptions are no longer valid is that we have been considering collections of sequences of (single agent, one-step) contractions that are organized in a tree as in Figure 3 above.

In our ongoing work we consider alternative models for interleaved contractions called *entrenchment systems* that are based on compositions of single agent entrenchment relations. One can prove a representation result to the effect that every selection system can be represented as an entrenchment system, and vice versa.

Our future work revolves around the idea of using other models of concurrency than interleaving, and determining the effects this has on the postulates describing multi-agent theory change.

Acknowledgments. Wiebe van der Hoek was supported in part by ESPRIT III BRWG project No. 8319 'ModelAge'.

A Proofs

Below we give proofs for results that were stated without proofs in earlier sections.

Theorem 3. *A contraction function $-$ satisfies the postulates of Definition 1 iff there exists a two-placed selection function s such that, for any theory T and formula ϕ,*

$$T - \phi = \bigcap s(T, (T \perp \phi)).$$

Proof. Let us first assume that $-$ is defined using a selection function s: we show that $-$ satisfies the required postulates. The Closure and Inclusion conditions follow immediately from the definition of s and the fact that theories are closed under intersection. Uniformity follows because if $T' \vdash \phi \leftrightarrow \psi$ for all subtheories $T' \subseteq T$, then $(T \perp \phi) = (T \perp \psi)$ and hence

$$s(T, (T \perp \phi)) = s(T, (T \perp \psi)).$$

Success is also clear: if $\nvdash \phi$, then $\phi \notin X$ for any $X \in (T \perp \phi)$, so $\phi \notin \bigcap s(T, T \perp \phi))$. As to the Relevance postulate, suppose $\psi \in T \setminus \bigcap s(T, (T \perp \phi))$. Then, there must be a $T' \in s(T, (T \perp \phi))$ with $\psi \notin T'$. By the definition of s, $T' \subseteq T$ and, by the definition of $(T \perp \phi)$ we must have $T', \psi \vdash \phi$.

Conversely, let $-$ satisfy the postulates of Definition 1. For any theories T and T' and formula ϕ such that $T - \phi = T'$, we have to guarantee that $\bigcap s(T, (T \perp \phi)) = T'$. We define $s(T, \Theta)$, with $\Theta \in 2^{\mathcal{T}}$ as follows.

$$s(T, \Theta) := \begin{cases} \{T\}, & \text{if } \Theta = \emptyset \\ \{S \in \Theta \mid T - \theta \subseteq S\}, & \text{if } \Theta = (T \perp \theta) \neq \emptyset \text{ for some formula } \theta \\ \Theta, & \text{otherwise.} \end{cases}$$

To see that s is a selection function, we first observe that $s(T, \emptyset) = \{T\}$ and if $\Theta \neq \emptyset$, then $s(T, \Theta)$ is a non-empty subset of Θ. It is also a function: suppose $T_1 = T_2$ and $\Theta_1 = \Theta_2$. If (T_1, Θ_1) is not a matching pair, then neither is (T_2, Θ_2) and we have

$$s(T_1, \Theta_1) = \Theta_1 = \Theta_2 = s(T_2, \Theta_2).$$

Otherwise, we may assume that $\Theta_1 = T_1 \perp \theta_1$ and $\Theta_2 = T_2 \perp \theta_2$ for some formulas θ_1 and θ_2. Thus $(T_1 \perp \theta_1) = (T_2 \perp \theta_2)$. Let T' be an arbitrary subtheory of $T_1 = T_2$. If $T' \vdash \theta_1$, we have $T' \notin (T_1 \perp \theta_1)$ and hence $T' \notin (T_2 \perp \theta_2)$. Thus, we have either that $T' \vdash \theta_2$ and then also $T' \vdash \theta_1 \leftrightarrow \theta_2$, or some $S \supseteq T'$ with $S \in (T_2 \perp \theta_2)$. The latter is impossible, since it would yield $S \in (T_1 \perp \theta_1)$ and $S \vdash \theta_1$. This proves that for any subtheory T' of T_1 we have $T' \vdash \theta_1 \leftrightarrow \theta_2$. By Uniformity and $T_1 = T_2$, we then have $T_1 - \theta_1 = T_2 - \theta_2$. From this we immediately get

$$\begin{aligned} s(T_1, \Theta_1) &= \{S \in \Theta_1 \mid T_1 - \theta_1 \subseteq S\} \\ &= \{S \in \Theta_2 \mid T_2 - \theta_2 \subseteq S\} \\ &= s(T_2, \Theta_2). \end{aligned}$$

Finally, we have to show that $\bigcap s(T, (T \perp \phi)) = T'$, whenever $T - \phi = T'$. We immediately have $T' \subseteq \bigcap s(T, (T \perp \phi))$. For the other inclusion, we first assume $\phi \in T$. Suppose we have some $\psi \notin T'$. By Relevance, we find an S' with $T' \subseteq S' \subseteq T$, for which $S', \psi \vdash \phi$ and $S' \nvdash \phi$. This S' can be expanded to an $S \supseteq S'$ such that $S \in (T \perp \phi)$ and $\psi \notin S$. We thus have

$$\psi \notin \bigcap s(T, (T - \phi)).$$

Finally, if $\phi \notin T$, then by Vacuity (which follows from Inclusion and Relevance), we have $T' = T$ and thus $\{T\} = s(T, (T \perp \phi))$, so that $\bigcap s(T, (T \perp \phi)) \subseteq T$. $\dashv$

Theorem 4. *If a contraction function $-$ satisfies the postulates of Definition 1 it also satisfies redundancy.*

Proof. Let T be a theory and suppose $\nvdash \phi$. Suppose furthermore that Z is a set of formulas, satisfying: (i) $T \cup Z$ is a theory, and (ii) for all $\zeta \in Z$: $\vdash \zeta \to \phi$. We have to prove: $T - \phi = (T \cup Z) - \phi$. If $Z \subseteq T$ the equation holds trivially, so let us assume the existence of a $\zeta \in Z \setminus T$. We now first show that $\phi \notin T$: if we would have $\phi \in T$, we reason as follows. Since $\zeta \in Z$, we have $(\zeta \vee \neg\phi) \in Z \cup T$. But $(\zeta \vee \neg\phi) \notin T$, since otherwise we would have, by $\phi, \phi \to \zeta \in T$ that $\zeta \in T$. Thus, $(\zeta \vee \neg\phi) \in Z$. By definition of Z, we have $\vdash (\zeta \vee \neg\phi) \to \phi$. Since $\vdash ((\zeta \vee \neg\phi) \to \phi) \to \phi)$, we would have $\vdash \phi$, contradicting one of the premises. Thus, $\phi \notin T$.

Now we can prove that $(T \cup Z) - \phi = T$. For $\subseteq$, suppose $\psi \in (T \cup Z) - \phi$. By Inclusion, we have $\psi \in T \cup Z$. If ψ would be in Z, we would have $\vdash \psi \to \phi$ and hence $\phi \in (T \cup Z) - \phi$, contradicting Success. Thus, $\psi \in T$. To see that also $T \subseteq (T \cup Z) - \phi$, let $\psi \in T$. Then $\psi \in T \cup Z$. If $\psi \notin (T \cup Z) - \phi$, by Relevance, we find a U with

$$(T \cup Z) - \phi \subseteq U \subseteq T \cup Z$$

and such that $U, \psi \vdash \phi$ and $U \nvdash \phi$. By the assumptions on Z, the latter implies that $U \subseteq T$. Since $\psi \in T$ and $U, \psi \vdash \phi$ we have $T \vdash \phi$ — a possibility we already excluded. Thus, $\psi \in (T \cup Z) - \phi$. $\dashv$

Theorem 5. *A contraction function $-$ satisfies the postulates of Definition 1 iff there exists a two-placed unified selection function s such that, for any theory T and formula ϕ,*

$$T - \phi = \bigcap s(T, (T \perp \phi)).$$

Proof. If s is a selection function, by Theorem 3 we find a contraction function $-$ satisfying the postulates of Definition 1. Conversely, suppose $-$ satisfies the postulates of Definition 1. By Proposition 4 we know that it also satisfies redundancy. We will show that the selection function s whose existence is guaranteed by Theorem 3, is unified. To do so, suppose

$$(U \perp \phi) = (V \perp \psi). \tag{4}$$

We have to show that $\bigcap s(U, (U \perp \phi)) = \bigcap s(V, (V \perp \psi))$. If $(U \perp \phi) = \emptyset$, the conclusion follows from the definition of s. So suppose $(U \perp \phi) \neq \emptyset$. We will first argue that

$$(U \perp \phi) = ((U \cap V) \perp \phi). \tag{5}$$

Suppose that $\chi \in U \setminus V$. Then $\chi \notin V$ and hence $\chi \notin Y$ for any $Y \in (V \perp \psi)$ and, by (4), $\chi \notin X$ for any $X \in (U \perp \phi)$. Since χ has been removed from all maximal subsets of U that do not entail ϕ, we must have $\vdash \chi \to \phi$. Thus

$$\alpha \in U \setminus V \Rightarrow \ \vdash \alpha \to \phi. \tag{6}$$

To prove the $\subseteq$-direction of (5), suppose $X \in (U \perp \phi)$. Then $X \nvdash \phi$ and by (6) we must also have $X \subseteq V$, and so $X \in ((U \cap V) \perp \phi)$. Conversely, suppose $X \in ((U \cap V) \perp \phi)$. Then $X \nvdash \phi$. Let χ be any formula in $U \setminus X$. If we can show that $X, \chi \vdash \phi$, we may conclude $X \in (U \perp \phi)$. Firstly, if $\chi \in V$, then, since $X \in ((U \cap V) \perp \phi)$, we have $X, \chi \vdash \phi$. If $\mu \notin V$ we have $\mu \in U \setminus V$, and by (6), $X, \mu \vdash \phi$. This proves (5), and, by a similar argument, we of course have $(V \perp \psi) = ((U \cap V) \perp \psi)$. Combining this with (4), we get

$$((U \cap V) \perp \phi) = ((U \cap V) \perp \psi). \tag{7}$$

Taking $T = U \cap V$ and $Z = U \setminus V$, Redundancy guarantees that $(U \cap V) - \phi = U - \phi$. Since $-$ is modeled by a selection function s, we have

$$\bigcap s\left(U \cap V, ((U \cap V) \perp \phi)\right) = \bigcap s\left(U, (U \perp \phi)\right).$$

Similarly, $\bigcap s(U \cap V, (U \cap V) \perp \psi) = \bigcap s(V, (V \perp \psi))$. From (7) we infer

$$\bigcap s(U \cap V, (U \cap V) \perp \phi) = \bigcap s(U \cap V, (U \cap V) \perp \psi),$$

so that we can finally conclude that $\bigcap s(U, (U \perp \phi)) = \bigcap s(V, (V \perp \psi))$, as required. $\dashv$

References

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

[2] M. Ben-Ari. *Principles of Concurrent and Distributed Programming.* Prentice Hall, 1990.

[3] A. Fuhrmann and S.-O. Hansson. A survey of multiple contractions. *Journal of Logic, Language and Information*, 3:39–75, 1994.

[4] P. Gärdenfors. *Knowledge in Flux.* The MIT Press, 1988.

[5] P. Gärdenfors and H. Rott. Belief revision. In D.M. Gabbay, C.J. Hogger, and J.A. Robinson, eds, *Handbook of Logic in AI and Logic Programming, Vol IV*, Oxford University Press, 1992.

[6] S.O. Hansson. Reversing the Levi identity. *Journal of Philosophical Logic*, 22:637–669, 1993.

[7] W. van der Hoek and M. de Rijke. Concurrent theory change. Manuscript, CWI, Amsterdam, 1995.

[8] N.E. Kfir-Dahav and M. Tennenholtz. Multi-agent belief revision. In *Proceedings TARK 96*, 1996, pages 175–194.

[9] H.F. Korth and A. Silberschwartz. *Database System Concepts.* McGraw-Hill, second edition, 1991.

[10] D. Lehmann. Belief revision, revised. In *Proceedings IJCAI 95*, 1995.

[11] S. Lindström and W. Rabinowicz. Epistemic entrenchment with incomparabilities and relational belief revision. In A. Fuhrmann and M. Morreau, eds, *The Logic of Theory Change*, Springer, 1991.

[12] D. Makinson. How to give it up: A survey of some formal aspects of the logic of theory change. *Synthese*, 62:347–363, 1985.